

MPLS
Internet-Draft
Intended status: Informational
Expires: January 30, 2020

Z. Zhang
Juniper Networks
S. Esale
Juniper Networks, Inc.
July 29, 2019

Resilient MPLS Rings and Multicast
draft-ietf-mpls-rmr-multicast-00

Abstract

With Resilient MPLS Rings (RMR), although all existing multicast procedures and solutions can work as is, there are optimizations that could be done for RSVP-TE P2MP tunnel signaling and Fast-ReRouting for both mLDP and RSVP-TE P2MP tunnels. This document describes RMR multicast on a high level, with detailed protocol procedure for RSVP-TE P2MP optimizations specified in a separate document. This document also discusses end to end multicast when there are RMRs.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on January 30, 2020.

Copyright Notice

Copyright (c) 2019 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

Internet-Draft

rmr-mcast

July 2019

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	P2MP/MP2MP Tunnels on a Ring	3
2.1.	Tunnel Protection and FRR	3
3.	End to End Tunnels with Rings	4
4.	End to End Native Multicast with Rings	4
4.1.	Native Multicast in the Global Routing Table	4
4.2.	mLDP Inband Signaling	4
4.3.	Overlay Multicast Services	4
4.3.1.	Tunnel Segmentation	5
5.	Summary	5
6.	Security Considerations	5
7.	IANA Considerations	5
8.	Acknowledgements	6
9.	References	6
9.1.	Normative References	6
9.2.	Informative References	6
	Authors' Addresses	7

[1.](#) Introduction

This document discusses how multicast works with Resilient MPLS Rings [[I-D.ietf-mpls-rmr](#)]. It is expected that readers are familiar with the concept and terms in [[I-D.ietf-mpls-rmr](#)].

All existing multicast procedures and solutions can work as is. This include both mpls multicast tunnels and end-to-end multicast that makes use of multicast tunnels. Ring topology is just a special case of general topologies so all existing RSVP-TE P2MP [[RFC4875](#)] and mLDP [[RFC6388](#)] tunnels can be set up using existing protocols and procedures. An Ingress Replication (IR) tunnel [[RFC7988](#)] consists a bunch of P2P LSPs, and it does not matter whether a component LSP is a plain old LSP or a Ring LSP.

On the other hand, there are optimizations that could be done for RSVP-TE P2MP tunnel signaling and Fast-ReRouting (FRR) for both mLDP and RSVP-TE P2MP tunnels. This document describes that on a high level, and discusses end to end multicast when there are RMRs even

though RMR could be transparent to multicast.

2. P2MP/MP2MP Tunnels on a Ring

Because mLDP label mapping messages are merged as they propagate from the leaves towards the root, ring topology does not lead to any further optimization in tunnel signaling.

However RSVP-TE P2MP tunnel signaling and procedures can be greatly optimized, as specified in [[I-D.zzhang-teas-rmr-rsvp-p2mp](#)].

2.1. Tunnel Protection and FRR

Each node on a ring signals two counter-rotating MP2P Ring LSPs to itself. As these LSPs are self-signaled after the discovery of the ring, they can be used to protect P2MP LSPs on ring. So neither mLDP nor RSVP-TE has to setup a separate P2P bypass LSPs for link and node protection. For instance, consider a ring with 8 nodes:




Figure 1: Ring with 8 nodes

Further, suppose a P2MP LSP is signaled with R0 as a root and R2, R4 and R6 as leaves. The P2MP LSP is formed as follows:





Figure 2: P2MP LSP

In the event of a link failure between R2 and R3, R2 the Point of Local Repair (PLR) tunnels P2MP LSP traffic on a anti-clockwise ring LSP to R3 the Merge Point (MP). Once the traffic is out of the ring

LSP on R3, it uses the regular P2MP LSP to reach R4. Similarly in the event of a node failure R3, R2 the PLR tunnels P2MP LSP traffic to R4 (the MP), which is also the leaf. Thus, the P2MP LSP uses the existing RSVP-TE ring LSPs for link and node protection.

3. End to End Tunnels with Rings

Consider a provider network that consists of one or more rings, optionally with a general topology connecting the rings. Multicast VPN [[RFC6514](#)], Ethernet VPN [[RFC7432](#)], VPLS [[RFC4761](#)] [[RFC4762](#)], or Global Table Multicast (GTM) via MVPN [[RFC7716](#)] overlay services are provided where end-to-end multipoint tunnels are needed across the entire network.

If the end to end tunnels are established by RSVP-TE P2MP, there is not much optimization that can be done for RMR, unless overlay-assisted tunnel segmentation is used. That is described in [Section 4.3.1](#).

If the end to end tunnels are established by mLDP and RSVP-TE signaling is desired on part of the network, mLDP Over Targeted Sessions [[RFC7060](#)] can be used (without the help from the overlay service) to stack part of an mLDP tunnel over a RSVP-TE P2MP tunnel. If the RSVP-TE P2MP tunnel is over a ring, then the optimization described earlier can be used.

4. End to End Native Multicast with Rings

Consider a network that consists of some rings. In this network,

end-to-end native multicast can take various forms described below.

[4.1.](#) Native Multicast in the Global Routing Table

This is typically signaled by PIM [[RFC7761](#)] end to end. This works for any topology and RMR does not make any difference.

[4.2.](#) mLDP Inband Signaling

This is specified in [[RFC6826](#)] [[RFC7246](#)] [[RFC7438](#)]. When part of a native (s,g) or (*,g) multicast tree needs to go over an mLDP domain, an mLDP tunnel is created for each multicast tree for the domain. RMR does not make any differences here.

[4.3.](#) Overlay Multicast Services

Overlay multicast services provided by MVPN/GTM/EVPN/VPLS use overlay multicast signaling to signal customer multicast state and tunnel binding. PE-PE multipoint underlay tunnels are used to distribute

multicast packets among PEs. Any kind of tunnel can be used, whether the provider network has rings or not, with or without the RMR related optimizations ([Section 3](#)).

[4.3.1.](#) Tunnel Segmentation

The MVPN/GTM/EVPN/VPLS PEs could span across ASes or areas. When the PE-PE multipoint tunnels cannot be signaled across AS/area boundaries, segmentation procedures can be used, as specified in [[RFC6514](#), [RFC7024](#)] and [I-D.ietf-bess-evpn-bum-procedure-updates]. With the base MVPN/GTM/EVPN/VPLS procedures, PEs advertise I/S-PMSE A-D routes to signal traffic to tunnel binding, and the routes carry type and identification of multi-point tunnels used to carry corresponding traffic. With segmentation, the ASBRs/ABRs become segmentation points and they change the tunnel type/identification when they re-advertise the routes to the next AS/area. With this, each AS/area has its own tunnel of different type/identification, stitched together by the ASBRs/ABRs.

With segmentation, different RMRs could have their own tunnels, and RSVP-TE P2MP optimizations for RMRs could be applied. Notice that this is different from [Section 3](#) in that overlay signaling is

involved.

[5.](#) Summary

As described above, multicast in the presence of RMRs can work as is. RSVP-TE P2MP tunnel signaling can be optimized (to be specified separately). Tunnel protection/FRR can also be optimized for mLDP/ RSVP-TE P2MP tunnels.

[6.](#) Security Considerations

This is an informational document that describes how existing multicast protocols can be used with RMR, as well as possible RMR specific enhancements that will be specified separately. There are no security concerns to be discussed here, as they are already discussed in existing protocols or will be discussed in the specification for the enhancements.

[7.](#) IANA Considerations

This document does not request any allocations from IANA. The RFC Editor is requested to remove this section before publication.

[8.](#) Acknowledgements

The authors sincerely thank Loa Andersson for his careful review, comments and suggestions.

[9.](#) References

[9.1.](#) Normative References

[I-D.ietf-mpls-rmr]

Kompella, K. and L. Contreras, "Resilient MPLS Rings", [draft-ietf-mpls-rmr-11](#) (work in progress), June 2019.

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

DOI 10.17487/RFC2119, March 1997,
<<https://www.rfc-editor.org/info/rfc2119>>.

9.2. Informative References

- [I-D.zzhang-teas-rmr-rsvp-p2mp]
Zhang, Z., Deshmukh, A., and R. Singh, "RSVP-TE P2MP Tunnels on RMR", [draft-zzhang-teas-rmr-rsvp-p2mp-02](#) (work in progress), January 2019.
- [RFC4761] Kompella, K., Ed. and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", [RFC 4761](#), DOI 10.17487/RFC4761, January 2007, <<https://www.rfc-editor.org/info/rfc4761>>.
- [RFC4762] Lasserre, M., Ed. and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", [RFC 4762](#), DOI 10.17487/RFC4762, January 2007, <<https://www.rfc-editor.org/info/rfc4762>>.
- [RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol – Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", [RFC 4875](#), DOI 10.17487/RFC4875, May 2007, <<https://www.rfc-editor.org/info/rfc4875>>.
- [RFC6388] 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/info/rfc6388>>.

- [RFC7060] Napierala, M., Rosen, E., and IJ. Wijnands, "Using LDP Multipoint Extensions on Targeted LDP Sessions", [RFC 7060](#), DOI 10.17487/RFC7060, November 2013, <<https://www.rfc-editor.org/info/rfc7060>>.
- [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](https://www.rfc-editor.org/info/rfc7761), DOI 10.17487/RFC7761, March 2016, <<https://www.rfc-editor.org/info/rfc7761>>.

[RFC7988] Rosen, E., Ed., Subramanian, K., and Z. Zhang, "Ingress Replication Tunnels in Multicast VPN", [RFC 7988](https://www.rfc-editor.org/info/rfc7988), DOI 10.17487/RFC7988, October 2016, <<https://www.rfc-editor.org/info/rfc7988>>.

Authors' Addresses

Zhaohui Zhang
Juniper Networks

EMail: zzhang@juniper.net

Santosh Esale
Juniper Networks, Inc.

EMail: sesale@juniper.net