Workgroup: MPLS WG
Internet-Draft: draft-kompella-mpls-nffrr-04
Published: 20 October 2023
Intended Status: Standards Track
Expires: 22 April 2024
Authors: K. Kompella W. Lin
Juniper Networks Juniper Networks

No Further Fast Reroute

Abstract

There are several cases where, once Fast Reroute has taken place (for MPLS protection), a second fast reroute is undesirable, even detrimental. This memo gives several examples of this, and proposes a mechanism to prevent further fast reroutes.

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 22 April 2024.

Table of Contents

1. Introduction

MPLS Fast Reroute (FRR) [RFC4090] [RFC5286] [RFC7490] is a useful and widely deployed tool for minimizing packet loss in the case of a link or node failure. This has not only proven to be very effective, it is often the reason for using MPLS as a data plane. FRR works for a variety of control plane protocols, including LDP, RSVP-TE, and SPRING. Furthermore, FRR is often used to protect MPLS services such as IP VPN and EVPN.

Having said this, there are case where, once FRR has taken place, if the packet encounters a second failure, a second FRR is not helpful, perhaps even disruptive. For example, the packet may loop until TTL expires. This can lead to link congestion and further packet loss. Thus, the attempt to prevent a packet from being dropped may instead affect many other packets. Note that the “second” failure may simply be another manifestation of the same failure; see Figure 1.

This memo proposes a mechanism for preventing further FRR once in cases where such further protection may be harmful. Several examples where this is the case are demonstrated as motivation. A solution using special-purpose labels (SPLs) is then offered. Some mechanisms

for distributing the capability to avoid further fast reroutes are also discussed, although these may be better placed in other documents in other Working Groups.

1.1. Other Approaches

[ALDT] has a more elaborate mechanism for preventing loops due to multiple failures. This involves marking the nodes redirecting traffic in a header (either individually, or as node groups), and dropping the packet at a transit node if its ID is in the header.

1.2. Terminology

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.

2. Motivation

A few cases are given where “further fast reroute” is harmful. Some of the cases are for MPLS services; others for “plain” MPLS forwarding.

2.1. EVPN (VPN/VPLS) Active-active Multihoming

Consider the following topology for multihoming an Ethernet VPN (EVPN [RFC7432]) Customer Edge (CE) device for protection against the failure of a Provider Edge (PE) device or a PE-CE link. To do so, there is a backup MPLS path between PE2 and PE3 (denoted by the starred line).

Figure 1: EVPN Multihoming

Suppose (known unicast) traffic goes from CE1 to CE2. With active-active multihoming, this traffic will be load-balanced between PE2 (to CE2 via link link1) and PE3 (to CE2 via link2). If link1 were to fail, PE2 can still get traffic for CE2 by sending it over the backup path to PE3 (and similarly for PE3 if link2 fails).

However, suppose CE2 is down. PE2 will assume link1 is down and send traffic for CE2 to PE3 over the backup path. PE3 (which thinks that link2 is down; note that the single real failure of CE2 being down is manifested as separate failures to PE2 and PE3) will protect this “second” failure by sending traffic for CE2 over the backup path to PE2. Thus, traffic will ping-pong between PE2 and PE3 until TTL expires.

Thus, the attempt to protect traffic to CE2 may end up doing more harm than good, by congesting the backup path between PE2 and PE3 and by giving PE2 and PE3 useless work to do.

A similar topology can be used in EVPN-Etree [RFC8317], EVPN-VPWS [RFC8214], IP VPN [RFC4364] or VPLS [RFC4761] [RFC4762]. In all these cases, the same looping behavior would occur for unicast traffic if CE2 is down.

2.2. RMR Protection

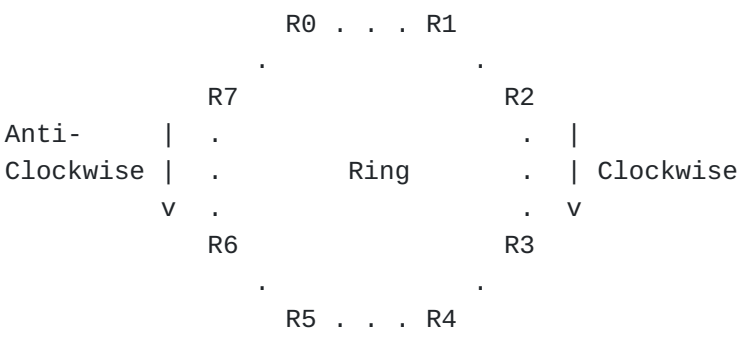

Figure 2: RMR Looping

In Resilient MPLS Rings (RMR), suppose traffic goes from a node, say R0, to a node, say R4, over a clockwise path. Protection consists of switching this traffic onto the anti-clockwise path to R4. This works well if a node or link between R0 or R4 is down. However, if node R4 itself is down, its adjacent neighbor R3, will send the traffic anti-clockwise to R4; when this traffic reaches R4’s other neighbor R5, it will return to N3, and so on, until TTL expires. [I-D.ietf-mpls-rmr] provides more details, and offers some means of mitigation. This memo offers a more elegant solution.

2.3. General MPLS forwarding

Consider the following topology:

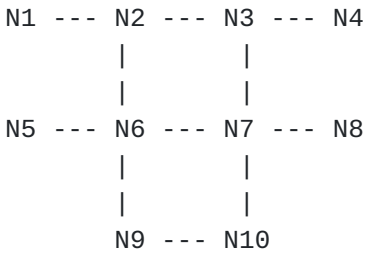

Figure 3: General MPLS Forwarding

Say link protection is configured for links N2-N3 and N6-N7. Link N2-N3 is protected by a bypass tunnel N2-N6-N7-N3, and link N7-N3 is protected by a bypass tunnel N7-N6-N2-N3. (These bypass tunnels may be set up using RSVP-TE [RFC3209] or via SPRING stacks [RFC8660].) Say furthermore that there is an LSP from N1 to N4 with path N1-N2-N3-N4, which asks for link protection. If link N2-N3 fails, traffic will take the path N1-N2-N6-N7-N3-N4.

Suppose, however, links N2-N3 and N7-N3 fail simultaneously. This may happen if they share fate (e.g., go over a common fiber conduit); it may also appear to happen if node N3 fails. Either way, first, the bypass protecting link N2-N3 kicks in, and traffic is sent to N3 via N6 and N7. However, when the traffic hits N7, the bypass for N7-N3 kicks in, and traffic is sent back to N2. Thus the traffic will loop between N2 and N7 until TTL expires, in the process congesting links N2-N6 and N6-N7.

Now consider an LSP: N5-N6-N7-N8. The link N6-N7 may be protected by the bypass N6-N2-N3-N7 or by N6-N9-N10-N7, or by load-balancing between these two bypasses. If both links N2-N3 and N6-N7 fail, then traffic that is protected via bypass N6-N2-N3-N7 will ping-pong between N6 and N2 until TTL expires; traffic protected via bypass N6-N9-N10-N7 will successfully make it to N8. If link N6-N7 is protected by load-balancing across the two bypass paths, then about half the traffic will loop between N6 and N2, and the rest will make it to N8.

While the above description is for protection using a bypass tunnel, the same principle applies to protection using Loop-Free Alternates [RFC5286] [RFC7490] or any of its variants (such as Topology Independent LFA).

3. Solution

To address this issue, we suggest the use of a SPL [RFC7274] called NFFRR (value TBD; suggested: 8). An alternate would be to use an extended SPL, whereby a pair of labels indicates that no further fast route is desired. However, in the case of SPRING MPLS bypass tunnels (Section 3.2.1) of depth N, this would triple the label

stack size. Using regular SPLs instead would only double the stack size.

To achieve loop-free fast rerouting for MPLS-based VPN networks, such as L3VPN, EVPN, VPWS, we can also use a different label for fast-rerouted data packets, separate from the label used for regular data packets. This fast reroute label allows the egress routers to distinguish fast-rerouted data packets from regular data packets, thus enabling loop-free reroute in response to another link or node failure.

It's important to note that the fast reroute label also functions as a service label for a VPN. Therefore for egress link protection, a dedicated reroute label would be required for each multihomed CE. For example, in the case of EVPN multihoming, each EVPN instance typically requires a distinct fast reroute label for each multihomed Ethernet Segment or virtual Ethernet segment. As the number of multihomed Ethernet segments, virtual Ethernet segments, and EVPN instances increases, so does the number of required fast reroute labels.

In comparison, the utilization of dedicated NFFRR to achieve loop-free fast rerouting offers several advantages:

1. Resource Efficiency: By using a dedicated NFFRR label unaffected by label allocation schemes or the number of multihomed CEs or VPNs, the number of label consumption is minimized. The NFFRR approach also reduces the need for additional routes and next-hop resources that would be otherwise required to forward fast-rerouted data packets based on individual fast reroute label.

2. Control Plane Efficiency: The use of a dedicated NFFRR label improves efficiency in the control plane. There's no need to signal individual fast reroute labels at the egress and then process them at the ingress. Routers supporting NFFRR recognize that no further routing is necessary based on the SPL label alone.

3. Uniform Approach: The same NFFRR label and scheme can be used for the different VPN services, such as L3VPN, EVPN, VPWS, etc., to achieve loop free egress link prorection. Furthermore, the use of the same NFFRR label can be extended beyond MPLS-based VPNs to other protocols offering bypass or fast reroute mechanisms in various network topologies, ensuring a uniform approach to achieving loop-free fast rerouting.

3.1. NFFRR for MPLS forwarding

To illustrate, we’ll first take the example of Figure 3, with MPLS paths signaled using RSVP-TE. This method can be used for paths that use SPRING stacks, but this will be detailed in a later version.

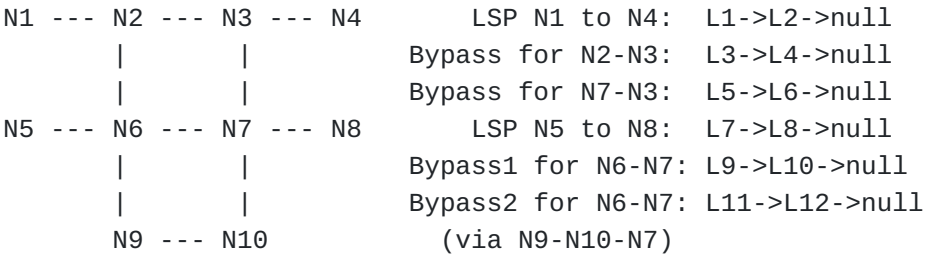

Figure 4: Example Using RSVP-TE LSPs

Node	Action	Next	New Pkt	Comment
N1	push L1	N2	[L1] pkt	ingress
N2	L1 -> L2	N3	[L2] pkt	
N3	pop L2	N4	pkt	PHP
N4	fwd pkt	–	–	continue

Table 1: Forwarding from N1 to N4

Note 1: “[L1 …]” denotes the label stack on the packet; pkt is the original packet received at ingress. “L1 -> L2” means swap label L1 with L2. “pop L2” means pop the top label L2. “fwd pkt” means forward the packet as usual.

Node	Action	Next	New Pkt	Comment
N2	push L3	N6	[L3] pkt	ingress
N6	L3 -> L4	N7	[L4] pkt	
N7	pop L4	N3	pkt	PHP

Table 2: Forwarding over the bypass for link N2-N3

Node	Action	Next	New Pkt	Comment
N7	push L5	N6	[L5] pkt	ingress
N6	L5 -> L6	N2	[L6] pkt	
N2	pop L6	N3	pkt	PHP

Table 3: Forwarding over Bypass1 for link N7-N3

Node	Action	Next	New Pkt	Comment
N1	push L1	N2	[L1] pkt	ingress
N2	L1 -> L2	N3	[L2] pkt	N3 X
N2	push L3	N6	[L3 L2] pkt	PLR
N6	L3 -> L4	N7	[L4 L2] pkt	
N7	pop L4	N3	[L2] pkt	merge
N3	pop L2	N4	pkt	PHP
N4	fwd pkt	–	–	continue

Table 4: Forwarding from N1 to N4 if link N2-N3 fails

Table 4 is obtained by composing Table 1 and Table 2.

Note 2: “N3 X” means “next hop N3 unavailable (because link N2-N3 failed)”.

Node	Action	Next	New Pkt	Comment
N1	push L1	N2	[L1] pkt	ingress
N2	L1 -> L2	N3	[L2] pkt	N3 X
N2	push L3	N6	[L3 L2] pkt	PLR
N6	L3 -> L4	N7	[L4 L2] pkt	
N7	pop L4	N3	[L2] pkt	N3 X’
N7	push L5	N6	[L5 L2] pkt	
N6	L5 -> L6	N2	[L6 L2] pkt	PLR
N2	pop L6	N3	[L2] pkt	N3 X
N2	push L3	N6	[L3 L2]	PLR
etc				loop!

Table 5: Forwarding from N1 to N4 if links N2-N3 and N7-N3 fail

Table 5 is obtained by composing Table 1, Table 2 and Table 3.

Note 3: “N3 X’” means “next hop N3 unavailable because link N7-N3 is down.

Note 4: While the impact of a loop is pretty bad, the impact of an ever-growing label stack (not illustrated here) and possible associated fragmentation on transit nodes may be worse.

3.2. Proposal

An LSR (typically a PLR) that wishes to prevent further FRRs after the first one can push an SPL, namely NFFRR, onto the label stack as follows:

Node	Action	Next	New Pkt	Comment
N1	push L1	N2	[L1] pkt	ingress
N2	L1 -> L2	N3	[L2] pkt	N3 X
N2	push L3, NFFRR	N6	[L3 NFFRR L2] pkt	PLR
N6	L3 -> L4	N7	[L4 NFFRR L2] pkt	
N7	pop L4, NFFRR	N3	[L2] pkt	merge
N3	pop L2	N4	pkt	PHP
N4	fwd pkt	–	–	continue

Table 6: Forwarding from N1 to N4 if link N2-N3 fails with NFFRR

Note 5: N2 can insert an NFFRR label only if it knows that all LSRs in the path can process it correctly. See Section 4 for some details on how this capability is communicated.

Node	Action	Next	New Pkt	Comment
N1	push L1	N2	[L1] pkt	ingress
N2	L1 -> L2	N3	[L2] pkt	N3 X
N2	push L3, NFFRR	N6	[L3 NFFRR L2] pkt	PLR
N6	L3 -> L4	N7	[L4 NFFRR L2] pkt	
N7	pop L4	N3	[NFFRR L2] pkt	N3 X
N7	check NFFRR	–	–	drop pkt

Table 7: Forwarding from N1 to N4 if links N2-N3 and N7-N3 fail with NFFRR

Note 6: “check NFFRR” means that, before N7 applies FRR (because link N7-N3 is down), N7 checks the label below the top label (or in this case, because of PHP, the top label itself). If this is the NFFRR label, N7 drops the packet rather than apply FRR.

3.2.1. NFFRR and SPRING

Suppose that, to protect link N2-N3, a bypass tunnel N2-N6-N7-N3 were instantiated using SPRING MPLS [RFC8660], in particular, using adjacency SIDs. If the corresponding labels for links N6-N7 and N7-N3 were L20 and L21, the bypass would consist of pushing the label stack [L20 L21] onto the packet and sending the packet to N6. To indicate that FRR has already occurred and to drop the packet rather than to try to protect the packet again, N2 would have to push [L20 NFFRR L21 NFFRR] onto the packet before sending it to N6. If the packet came from N1 with label L1, N2 would send a packet with label stack [L20 NFFRR L21 NFFRR L2] to N6.

N6 would see L20, pop it, note the NFFRR label and pop it, then attempt to send the packet to N7. If the link N6-N7 is down, N6 drops the packet. Otherwise, N7 gets the packet, sees L21, pops it, sees NFFRR, pops it and tries to send the packet to N3. If link N7-N3 is down, N7 drops the packet. Otherwise, N3 gets the packet with L2, swaps with with L3 and sends it to N4.

Note that with SPRING MPLS, the NFFRR label needs to be repeated for each label in the bypass stack. Hence the request for a “regular” SPL rather than an extended SPL.

3.3. NFFRR for MPLS Services

First, we illustrate known unicast EVPN forwarding:

Node	Action	Next	Packet	Comment
PE1	send to CE2	PE2	[T1 S2] pkt	EVPN
PE2	send to CE2	link1	pkt	done!

Table 8

Note: T1/T2/T3 are the transport labels for PE1/PE3/PE2 to reach PE2/PE2/PE3 respectively. S2/S3 are the service labels announced by PE2/PE3 for CE2.

Then, we show what happens when CE2 is down without NFFRR:

Node	Action	Next	Packet	Comment
PE1	send to CE2	PE2	[T1 S2] pkt	EVPN
PE2	send to CE2	link1	—	link1 X
PE2	send to CE2	PE3	[T3 S3] pkt	eFRR
PE3	send to CE2	link2	—	link2 X
PE3	send to CE2	PE2	[T2 S2] pkt	eFRR
PE2	send to CE2	link1	—	link1 X
PE2	send to CE2	PE3	[T3 S3] pkt	eFRR
…				loop!

Table 9

Note: link1/link2 X means link1/link2 is down. eFRR refers to EVPN multihoming FRR.

In the case of MPLS services such as EVPN Figure 1, the NFFRR label is inserted below the service label, as shown below:

Node	Action	Next	Packet	Comment
PE1	send to CE2	PE2	[T1 S2] pkt	EVPN
PE2	send to CE2	link1	—	link1 X
PE2	send to CE2	PE3	[T3 S2 NFFRR] pkt	eFRR
PE3	send to CE2	link2	—	link2 X
PE3	drop pkt	—	—	check NFFRR

Table 10

Note: “check NFFRR” is as above.

3.4. NFFRR for RMR

As described in Figure 2, packets will loop until TTL expires if the destination node in an RMR ring (here, R4) fails. The solution in this case is that the first node to apply RMR protection (R3) pops the current RMR transport label being used, sees that the next label is not NFFRR (so protection is allowed), pushes an NFFRR label and then the RMR transport label for the reverse direction.

When R5 receives the packet, it sees that the next link is down, pops the RMR transport label, sees the NFFRR label and drops the packet. Thus, the loop is avoided.

4. Signaling NFFRR Capability

4.1. Signaling NFFRR Capability for MPLS Services with BGP

The ideal choice would be an attribute consisting of a bit vector of node capabilities, one bit of which would be the capability of processing the NFFRR SPL below the BGP service label. This would be used by BGP L2VPN, BGP VPLS, EVPN, E-Tree and E-VPWS. An alternative is to use the BGP Capabilities Optional Parameter [I-D.ietf-idr-next-hop-capability]. Details to be worked out.

4.2. Signaling NFFRR Capability for MPLS Services with Targeted LDP

One approach to signaling NFFRR capability for MPLS services signaled with targeted LDP is to introduce a new LDP TLV called the NFFRR Capability TLV as an Optional Parameter in the Label Mapping Message [RFC5036]. This TLV has Type TBD (suggested: 0x0207) and Length 0.

Another approach is to use LDP Capabilities [RFC5561]; this approach has the advantage that it deals with capabilities on a node basis rather than on a per label mapping basis. However, there don’t appear to be other documents using this approach.

4.3. Signaling NFFRR Capability for MPLS Forwarding

The authors suggest signaling a router’s ability to process the NFFRR SPL using the Link State Router TE Node Capabilities [RFC5073], which works for both IS-IS and OSPF. A new TE Node Capability bit, the N bit (suggested value 5) indicates that the advertising node is capable of processing the NFFRR SPL.

5. IANA Considerations

If this draft is deemed useful, a way to signal that No Further Fast-route should be performed on a packet will be needed. There are two approaches: allocate an SPL for NFFRR: if so, we suggest the early allocation of label 8 for this. Alternatively, if [I-D.kompella-mpls-mspl4fa] (or similar) is adopted, allocate a forwarding action bit saying whether or not to do FRR.

Furthermore, means of signaling the ability to process the NFFRR SPL/bit should be defined for IS-IS, OSPF, LDP and BGP.

The following update is suggested for the Link State Router TE Node Capabilities registry:

Bit	Name	Reference
5	NFFRR	This document

Table 11

The following update is suggested for the TLV Type Name Space of the Label Distribution Protocol (LDP) Parameters registry:

Type	Name	Reference
0x0207	NFFRR	This document

Table 12

6. Security Considerations

A malicious or compromised LSR can insert NFFRR into a label stack, preventing FRR from occurring. If so, protection will not kick in for failures that could have been protected, and there will be unnecessary packet loss.

7. References

7.1. Normative References

[I-D.kompella-mpls-mspl4fa] Kompella, K., Beeram, V. P., Saad, T., and I. Meilik, "Multi-purpose Special Purpose Label for Forwarding Actions", Work in Progress, Internet-Draft, draft-kompella-mpls-mspl4fa-03, 10 July 2022, <https://datatracker.ietf.org/doc/html/draft-kompella-mpls-mspl4fa-03>.

[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>.

[RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, DOI 10.17487/RFC5036, October 2007, <https://www.rfc-editor.org/info/rfc5036>.

[RFC5073] Vasseur, J.P., Ed. and J.L. Le Roux, Ed., "IGP Routing Protocol Extensions for Discovery of Traffic Engineering Node Capabilities", RFC 5073, DOI 10.17487/RFC5073, December 2007, <https://www.rfc-editor.org/info/rfc5073>.

[RFC7274] Kompella, K., Andersson, L., and A. Farrel, "Allocating and Retiring Special-Purpose MPLS Labels", RFC 7274, DOI

10.17487/RFC7274, June 2014, <https://www.rfc-editor.org/info/rfc7274>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

7.2. Informative References

[ALDT] Merling, D., Braun, W., and M. Menth, "Efficient Data Plane Protection for SDN", June 2018, <https://atlas.informatik.uni-tuebingen.de/~menth/papers/Menth18g.pdf>.

[I-D.ietf-idr-next-hop-capability] Decraene, B., Kompella, K., and W. Henderickx, "BGP Next-Hop dependent capabilities", Work in Progress, Internet-Draft, draft-ietf-idr-next-hop-capability-08, 8 June 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-next-hop-capability-08>.

[I-D.ietf-mpls-rmr] Kompella, K. and L. M. Contreras, "Resilient MPLS Rings", Work in Progress, Internet-Draft, draft-ietf-mpls-rmr-14, 14 February 2021, <https://datatracker.ietf.org/doc/html/draft-ietf-mpls-rmr-14>.

[RFC3209] 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/info/rfc3209>.

[RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, DOI 10.17487/RFC4090, May 2005, <https://www.rfc-editor.org/info/rfc4090>.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <https://www.rfc-editor.org/info/rfc4364>.

[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>.

[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>.

[RFC5561] Thomas, B., Raza, K., Aggarwal, S., Aggarwal, R., and JL. Le Roux, "LDP Capabilities", RFC 5561, DOI 10.17487/RFC5561, July 2009, <https://www.rfc-editor.org/info/rfc5561>.

[RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <https://www.rfc-editor.org/info/rfc7432>.

[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>.

[RFC8214] Boutros, S., Sajassi, A., Salam, S., Drake, J., and J. Rabadan, "Virtual Private Wire Service Support in Ethernet VPN", RFC 8214, DOI 10.17487/RFC8214, August 2017, <https://www.rfc-editor.org/info/rfc8214>.

[RFC8317] Sajassi, A., Ed., Salam, S., Drake, J., Uttaro, J., Boutros, S., and J. Rabadan, "Ethernet-Tree (E-Tree) Support in Ethernet VPN (EVPN) and Provider Backbone Bridging EVPN (PBB-EVPN)", RFC 8317, DOI 10.17487/RFC8317, January 2018, <https://www.rfc-editor.org/info/rfc8317>.

[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>.

Authors' Addresses

Kireeti Kompella
Juniper Networks
1133 Innovation Way
Sunnyvale, CA 94089
United States

Email: kireeti.kompella@gmail.com

Wen Lin

Juniper Networks
1133 Innovation Way
Sunnyvale, CA 94089
United States

Email: wlin@juniper.net