

Workgroup: Network Working Group
Internet-Draft:
draft-chen-spring-srmppls-frr-ex-03
Published: 18 October 2023
Intended Status: Standards Track
Expires: 20 April 2024
Authors: H. Chen Z. Hu A. Wang
 Futurewei Huawei Technologies China Telecom
 Y. Liu G. Mishra
 China Mobile Verizon Inc.

SR-MPLS FRR Extension

Abstract

The current SR FRR such as TI-LFA provides fast re-route protection for the failure of a node on an SR-MPLS path by the neighbor upstream node as point of local repair (PLR) of the failed node. However, once the IGP converges, the SR FRR is no longer sufficient to forward traffic of the path around the failure, since the non-neighbor upstream node of the failed node will no longer have a route to the failed node. This document describes a simple mechanism to extend the fast re-route protection for the failure on an SR-MPLS path after the IGP converges. The mechanism protects the node SID, adjacency SID and binding SID of the failed node on the path.

Requirements Language

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

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 20 April 2024.

Copyright Notice

Copyright (c) 2023 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. Example SR-MPLS FRR Extension](#)
 - [2.1. SR-MPLS Path with no BSID](#)
 - [2.1.1. Without any Failure](#)
 - [2.1.2. Before IGP Converges on Failure](#)
 - [2.1.3. After IGP Converges on Failure](#)
 - [2.2. SR-MPLS Path with BSID](#)
 - [2.2.1. Without any Failure](#)
 - [2.2.2. Before IGP Converges on Failure](#)
 - [2.2.3. After IGP Converges on Failure](#)
- [3. Procedures](#)
 - [3.1. Procedure on Non-neighbor Upstream Node](#)
 - [3.2. Procedure on Neighbor Upstream Node](#)
- [4. Security Considerations](#)
- [5. IANA Considerations](#)
- [6. References](#)
 - [6.1. Normative References](#)
 - [6.2. Informative References](#)
- [Acknowledgments](#)
- [Authors' Addresses](#)

1. Introduction

[[I-D.ietf-rtgwg-segment-routing-ti-lfa](#)] describes an SR FRR mechanism that provides fast re-route protection for the failure of a node on an SR-TE path by the neighbor upstream node as point of local repair (PLR) of the failed node. However, once the IGP converges, the SR FRR is no longer sufficient to forward traffic of the path around the failure, since the non-neighbor upstream node of the failed node will no longer have a route to the failed node and drop the traffic.


[[I-D.ietf-spring-segment-protection-sr-te-paths](#)] proposes a solution in which a hold-down timer is configured on every node in a network. After the IGP converges on a node failure, when a node is going to delete the route to the failed node, instead of programming a route delete, it programs a tunnel/path to the node consisting of the Node SID of the nearside neighbor of the failed node followed by the original path in the packet. The modified path will be in force until the hold-down timer expires.

This document describes a simple mechanism to extend the fast re-route protection for the failure on an SR-MPLS path after the IGP converges. The mechanism protects the node SID, adjacency SID and binding SID of the failed node on the path.

2. Example SR-MPLS FRR Extension

This section illustrates the extension to SR-MPLS FRR for the failure on SR-MPLS paths after the IGP converges through examples. It shows the procedure on every related node on each path without any failure, with a failure before and after the IGP converges on the failure.

[Figure 1](#) shows an example topology with two SR-MPLS paths: Path 1 and Path 2 for explaining the extension. They go through the same nodes and links, but represented differently. The former does not have any binding SID (BSID). The latter has a BSID of node N (BSID-N).



Path 1:A->P1->N->Q1->C, indicated at A by {SID-P1,SID-N,SID-Q1,SID-C}

Path 2:A->P1->N->Q1->C, indicated at A by {SID-P1,SID-N,BSID-N} and
BSID-N associated with SID list {SID-Q1,SID-C}

Figure 1: Example Topology with SR-MPLS Paths

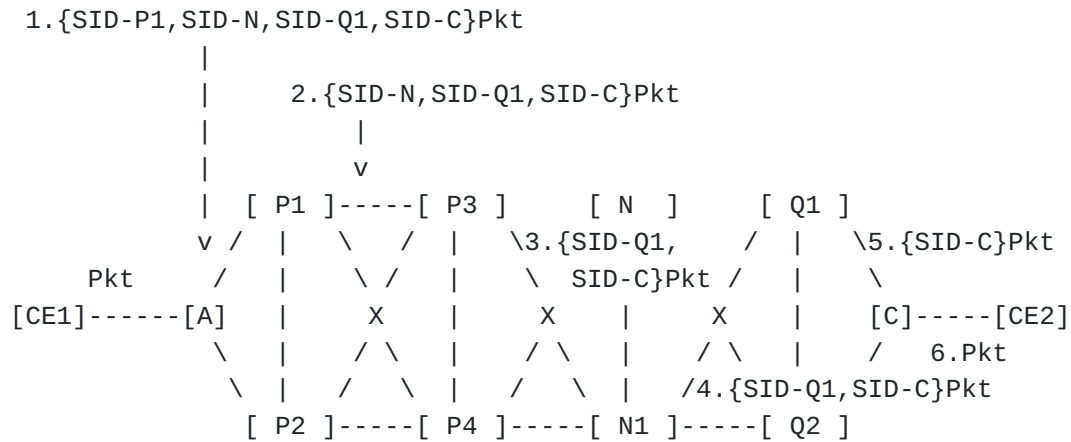
SR-MPLS Path 1: A->P1->N->Q1->C is indicated at node A by node SID of P1 (SID-P1), node SID of N (SID-N), node SID of Q1 (SID-Q1) and node SID of C (SID-C). SR-MPLS Path 2 is indicated at node A by {SID-P1,SID-N,BSID-N}, where BSID-N (binding SID of N) is associated with SID list {SID-Q1,SID-C}.

6.

Node C pops its SID-C and gets the packet without any SIDs, which is represented by "6.Pkt".

2.1.2. Before IGP Converges on Failure

Figure 3 shows the result of executing procedure on each related node of SR-MPLS path 1 with no BSID when node N failed and before the IGP converges on the failure.



Path 1: A->P1->N->Q1->C, indicated at A by {SID-P1, SID-N, SID-Q1, SID-C}

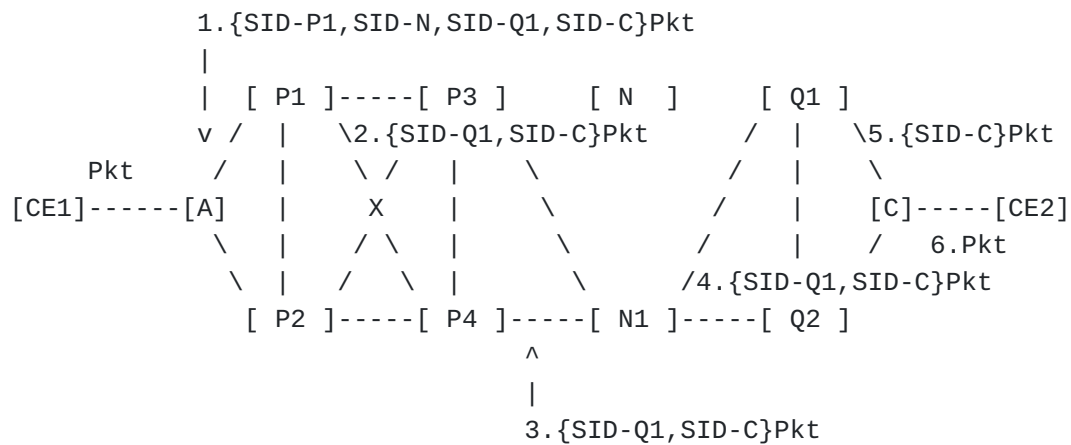
Figure 3: Path with no BSID when N failed and before IGP converges

The results from nodes A, P1, Q1 and C are the same as those described in Section 2.1.1. The results from neighbor upstream node P3 of N and node N1 are as follows.

3. After detecting the failure of N, the neighbor upstream node P3 of N pops SID-N from the packet received, re-routes the packet to node Q1 via node N1 without going through failed N. The packet sent to N1 is represented by "3. {SID-Q1, SID-C}Pkt".
4. Node N1 sends the packet to Q1 according to the top SID (SID-Q1) in the packet. The packet sent to Q1 is represented by "4. {SID-Q1, SID-C}Pkt".

2.1.3. After IGP Converges on Failure

Figure 4 shows the result of executing procedure on each related node on SR-MPLS path 1 with no BSID when node N failed and after the IGP converges on the failure.



Path 1:A->P1->N->Q1->C, indicated at A by {SID-P1,SID-N,SID-Q1,SID-C}

Figure 4: Path with no BSID when N failed and after IGP converges

The results from nodes A, N1, Q1 and C are the same as those described above in [Section 2.1.2](#). The results from nodes P1 and P4 are as follows.

2. Since SID-N is a failed node SID, non-neighbor upstream node P1 pops SID-N from the packet, and sends the packet to P4 according to the top SID (SID-Q1) in the packet along the IGP shortest path to Q1. The packet sent to P4 is represented by "2.{SID-Q1,SID-C}Pkt".
3. Node P4 sends the packet to N1 according to the top SID (SID-Q1) in the packet received. The packet sent to N1 is represented by "3.{SID-Q1,SID-C}Pkt".

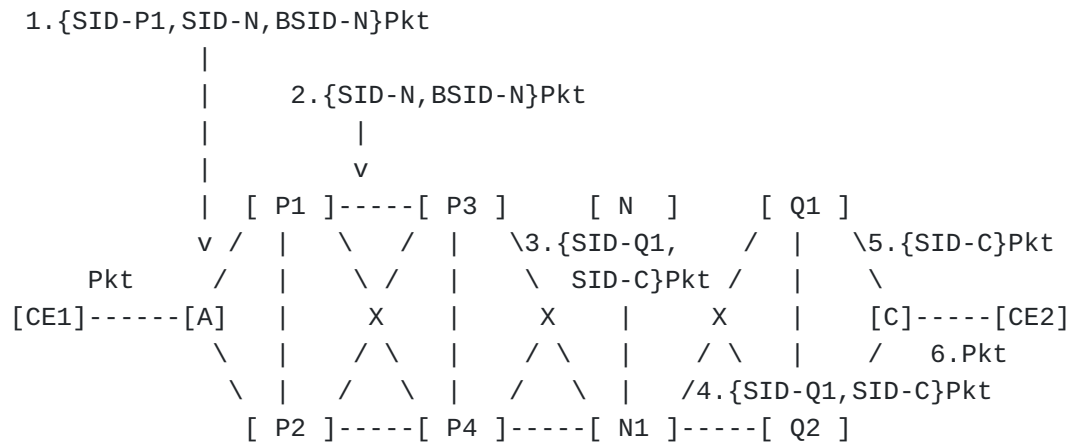
The non-neighbor upstream node P1 may determines whether a SID is a failed SID in the following way:

IF there is a RIB/FIB entry for the SID (e.g., SID-N) and
 then the entry for the SID is to be removed after a SPF
 THEN the SID is a failed SID.

2.2. SR-MPLS Path with BSID

2.2.1. Without any Failure

[Figure 5](#) shows the result of executing procedure on each related node of SR-MPLS path 2 with BSID-N without any failure.



Path 2: A->P1->N->Q1->C, indicated at A by {SID-P1, SID-N, BSID-N} and BSID-N associated with SID list {SID-Q1, SID-C}

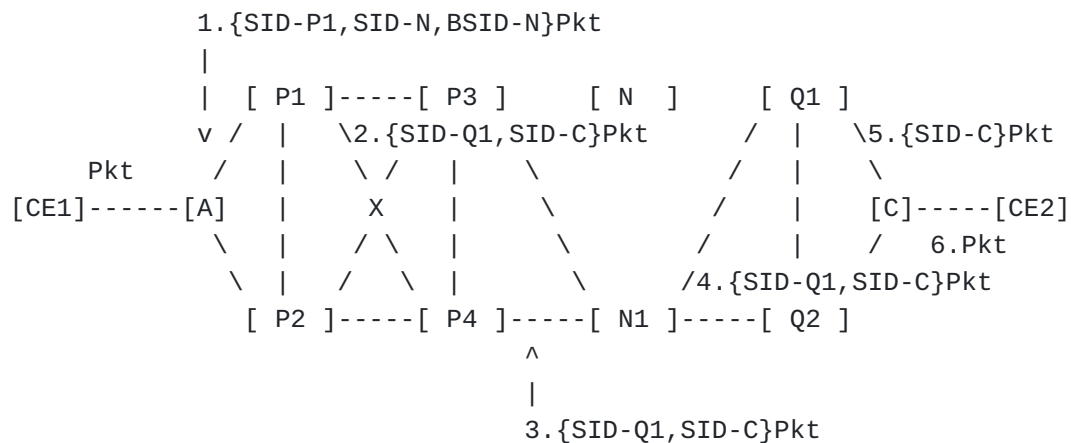
Figure 6: Path with BSID-N when N failed and before IGP converges

The results from nodes A, P1, Q1 and C are the same as those described in [Section 2.2.1](#). The results from neighbor upstream node P3 of N and node N1 are as follows.

3. After detecting the failure of N, the neighbor upstream node P3 of N pops SID-N from the packet received, replaces BSID-N in the packet with SID list {SID-Q1, SID-C}, and re-routes the packet to node Q1 via node N1 without going through failed N. The packet sent to N1 is represented by "3.{SID-Q1, SID-C}Pkt".
4. Node N1 sends the packet to Q1 according to the top SID (SID-Q1) in the packet. The packet sent to Q1 is represented by "4.{SID-Q1, SID-C}Pkt".

2.2.3. After IGP Converges on Failure

[Figure 7](#) shows the result of executing procedure on each related node on SR-MPLS path 2 with BSID-N when node N failed and after the IGP converges on the failure.



Path 2: A->P1->N->Q1->C, indicated at A by {SID-P1, SID-N, BSID-N} and BSID-N associated with SID list {SID-Q1, SID-C}

Figure 7: Path with BSID-N when N failed and after IGP converges

The results from nodes A, Q1 and C are the same as those described above in [Section 2.2.2](#). The results from nodes P4 and N1 are the same as those described above in [Section 2.1.3](#). The result from node P1 is as follows.

2. Since SID-N is a failed node SID, non-neighbor upstream node P1 pops SID-N from the packet, replaces BSID-N in the packet with SID list {SID-Q1, SID-C} and sends the packet to P4 according to the top SID (SID-Q1) in the packet along the IGP shortest path to Q1. The packet sent to P4 is represented by "2. {SID-Q1, SID-C} Pkt".

3. Procedures

This section presents the procedures on a neighbor upstream node and a non-neighbor upstream node of node N on an SR MPLS path.

3.1. Procedure on Non-neighbor Upstream Node

For an SR-MPLS path with the node SID of node N (SID-N), suppose that node X is an non-neighbor upstream node of node N along the path, wherein SID-N follows the node SID of node X (SID-X) or the adjacency SID to node X in the packet to be transported by the path. For example, SR-MPLS Path 1 and Path 2 in [Figure 1](#) are from A to P1 to N to Q1 to C. Node P1 is the non-neighbor upstream node of N. At ingress node A of the path, node A adds the SIDs in the SID list into the packet to be transported by the path.

Without any failure or from the failure of node N to IGP convergence on the failure, the non-neighbor upstream node (such as P1) of node N pops its SID from the packet if any and sends the packet to the next hop node toward node N along the IGP shortest path to N.

After node N failed and from the IGP convergence on the failure to global reroute, the non-neighbor upstream node (such as P1) of node N pops its SID (such as SID-P1) from the packet if any, pops SID-N from the packet and does one of the following:

- a. If the current top SID in the packet is a node SID of a node named Nx, sends the packet toward Nx along the IGP shortest path to Nx.
- b. If the current top SID in the packet is an adjacency SID of node N, obtains the remote node of the adjacency from the adjacency SID, replaces the adjacency SID with the node SID of the remote node, and sends the packet toward the remote node along the IGP shortest path to the remote node.
- c. If the current top SID in the packet is a Binding SID (BSID) of node N, replaces the BSID in the packet with the SID list associated with the BSID, and does a. or b. according to the current top SID in the packet (i.e., does a. if it is a node SID; does b. if it is an adjacency SID of N). Note: Distributing the information about the BSID of N (including the BSID, the SID list, the ID of N) to upstream nodes of N is out of scope of this document, but described in [[I-D.chen-pce-mbinding](#)] and [[I-D.chen-idr-mbinding](#)].

3.2. Procedure on Neighbor Upstream Node

Suppose that node Y is the neighbor upstream node of node N on an SR-MPLS path. Without any failure, node Y sends the packet received from the path to node N according to the top SID in the packet. For example, SR-MPLS Path 1 and Path 2 in [Figure 1](#) are from A to P1 to N to Q1 to C. Node P3 is the neighbor upstream node of N suppose that the shortest path from P1 to N is from P1 to P3 to N. When node Y detects the failure of node N, node Y pops SID-N from the packet if SID-N is the top SID of the packet, and does one of the following steps (where steps a and b are TI-LFA in [[I-D.ietf-rtgwg-segment-routing-ti-lfa](#)]):

- a. If the current top SID in the packet is node SID of node Nx, sends the packet toward node Nx without going through failed N (i.e., {SID-Nx,...} in the packet would be {RL(Q),SID-Nx,...}, where RL(Q) is the repair list redirecting the packet to node Q, whose path to SID-Nx is not affected by the failure).
- b. If the current top SID in the packet is an adjacency SID of N, gets remote node R of the adjacency from the adjacency SID, replaces the top SID in the packet with node SID of R, and sends the packet to node R according to the top SID in the packet without going through failed N (i.e., {adj(N-R),...} in the

packet would be {RL(Q),SID-R,...}, where RL(Q) is the repair list redirecting the packet to node Q, whose path to SID-R is not affected by the failure).

- c. If the current top SID in the packet is a BSID of N, replaces the BSID with the SID list associated with the BSID, and does a. or b. according to the current top SID in the packet (i.e., does a. if it is a node SID; does b. if it is an adjacency SID of N).

4. Security Considerations

TBD.

5. IANA Considerations

No requirements for IANA.

6. References

6.1. Normative References

[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>>.

[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>>.

6.2. Informative References

[I-D.chen-idr-mbinding] Chen, H., Decraene, B., Mishra, G. S., Fan, Y., Wang, A., and X. Liu, "BGP for Mirror Binding", Work in Progress, Internet-Draft, draft-chen-idr-mbinding-02, 10 May 2023, <<https://datatracker.ietf.org/doc/html/draft-chen-idr-mbinding-02>>.

[I-D.chen-pce-mbinding] Chen, H., Decraene, B., Mishra, G. S., Wang, A., Liu, X., and L. Liu, "PCE for Mirror Binding", Work in Progress, Internet-Draft, draft-chen-pce-mbinding-02, 8 October 2023, <<https://datatracker.ietf.org/doc/html/draft-chen-pce-mbinding-02>>.

[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", Work in Progress, Internet-Draft, draft-ietf-rtgwg-segment-routing-ti-

lfa-11, 30 June 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-rtgwg-segment-routing-ti-lfa-11>>.

[I-D.ietf-spring-segment-protection-sr-te-paths]

Hegde, S., Bowers, C., Litkowski, S., Xu, X., and F. Xu, "Segment Protection for SR-TE Paths", Work in Progress, Internet-Draft, draft-ietf-spring-segment-protection-sr-te-paths-05, 27 September 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-segment-protection-sr-te-paths-05>>.

Acknowledgments

The authors would like to thank Joel Halpern, Andrew Stone, Yao Liu, and Jeff Tantsura for their comments to this work.

Authors' Addresses

Huaimo Chen
Futurewei
Boston, MA,
United States of America

Email: Huaimo.chen@futurewei.com

Zhibo Hu
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing
100095
China

Email: huzhibo@huawei.com

Aijun Wang
China Telecom
Beiqijia Town, Changping District
Beijing
102209
China

Email: wangaj3@chinatelecom.cn

Yisong
China Mobile
510000
China

Email: liuyisong@chinamobile.com

Gyan S. Mishra
Verizon Inc.
13101 Columbia Pike
Silver Spring, MD 20904
United States of America

Phone: [301 502-1347](tel:3015021347)

Email: gyan.s.mishra@verizon.com