

Workgroup: Network Working Group
Internet-Draft:
draft-ietf-rtgwg-srv6-egress-protection-09
Published: 13 March 2023
Intended Status: Standards Track
Expires: 14 September 2023
Authors: Z. Hu H. Chen H. Chen P. Wu
 Huawei Futurewei China Telecom Huawei
 M. Toy C. Cao T. He L. Liu
 Verizon China Unicom China Unicom Fujitsu
 X. Liu
 IBM Corporation

SRv6 Path Egress Protection

Abstract

This document describes protocol extensions for protecting the egress node of a Segment Routing for IPv6 (SRv6) path or tunnel.

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 [RFC 2119](#) [[RFC2119](#)].

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 14 September 2023.

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. Terminologies](#)
- [3. SR Path Egress Protection](#)
 - [3.1. Mechanism](#)
 - [3.2. Example](#)
- [4. Extensions to IGP for Egress Protection](#)
 - [4.1. Extensions to IS-IS](#)
 - [4.2. Extensions to OSPF](#)
- [5. Security Considerations](#)
- [6. IANA Considerations](#)
 - [6.1. SRv6 Endpoint Behaviors](#)
 - [6.2. IS-IS](#)
 - [6.3. OSPFv3](#)
- [7. References](#)
 - [7.1. Normative References](#)
 - [7.2. Informative References](#)
- [Acknowledgments](#)
- [Authors' Addresses](#)

1. Introduction

The fast protection of a transit node of a Segment Routing (SR) path or tunnel is described in [[I-D.ietf-rtgwg-segment-routing-ti-lfa](#)]. [[RFC8400](#)] specifies the fast protection of egress node(s) of an MPLS TE LSP tunnel including P2P TE LSP tunnel and P2MP TE LSP tunnel in details. However, these documents do not discuss the fast protection of the egress node of a Segment Routing for IPv6 (SRv6) path or tunnel.

This document fills that void and presents protocol extensions for the fast protection of the egress node of an SRv6 path or tunnel. Egress node and egress, fast protection and protection as well as SRv6 path and SRv6 tunnel will be used exchangeably below.

There are a number of topics related to the egress protection, which include the detection of egress node failure, the relation between egress protection and global repair, and so on. These are discussed in details in [[RFC8679](#)].

2. Terminologies

The following terminologies are used in this document.

SR: Segment Routing

SRv6: SR for IPv6

SRH: Segment Routing Header

SID: Segment Identifier

LSA: Link State Advertisement in OSPF

LSP: Label Switched Path in MPLS or Link State Protocol PDU in IS-IS

PDU: Protocol Data Unit

LS: Link State, which is LSA in OSPF or LSP in IS-IS

TE: Traffic Engineering

SA: Source Address

DA: Destination Address

P2MP: Point-to-MultiPoint

P2P: Point-to-Point

CE: Customer Edge

PE: Provider Edge

LFA: Loop-Free Alternate

TI-LFA: Topology Independent LFA

BFD: Bidirectional Forwarding Detection

VPN: Virtual Private Network

L3VPN: Layer 3 VPN

VRF: Virtual Routing and Forwarding

FIB: Forwarding Information Base

PLR: Point of Local Repair

BGP:

Border Gateway Protocol

IGP: Interior Gateway Protocol**OSPF:** Open Shortest Path First**IS-IS:** Intermediate System to Intermediate System**3. SR Path Egress Protection**

This section describes the mechanism of SR path egress protection and illustrates it through an example.

3.1. Mechanism

[Figure 1](#) is used to explain the mechanism of SR path egress node protection.

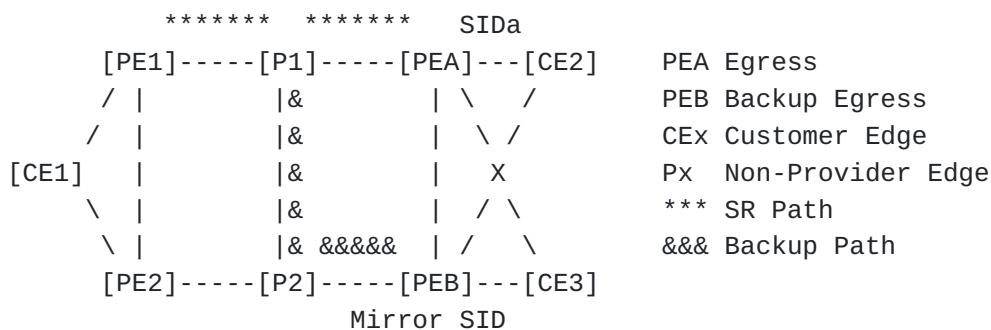


Figure 1: PEB Protects Egress PEA of SR Path

Where node PEA is the egress of the SR path from PE1 to PEA, and has SIDA which is the active segment in the packet from the SR path at PEA. Node PEB is the backup egress (or say protector) to provide the protection for egress (or say primary egress) PEA. Node P1 is the direct previous/upstream hop of egress PEA and acts as PLR (refer to [\[I-D.ietf-rtgwg-segment-routing-ti-lfa\]](#)) to support the protection for PEA.

When PEB is selected as a backup egress to protect the egress PEA, a Mirror SID (refer to Section 5.1 of [\[RFC8402\]](#)) is configured on PEB to protect PEA. PEB advertises this information through IGP, which includes the Mirror SID and the egress PEA. The information is represented by <PEB, PEA, Mirror SID>, which indicates that PEB protects PEA with Mirror SID.

After PEA receives the information <PEB, PEA, Mirror SID>, it may send the forwarding behavior of the SIDA at PEA to PEB with the Mirror SID using some protocols such as BGP if PEB can not obtain

this behavior from other approaches and PEB wants to protect SIDA of PEA. How to send the forwarding behavior of the SIDA to PEB is out scope of this document.

When PEB gets the forwarding behavior of the SIDA of PEA from PEA or other means, it adds a forwarding entry for the SIDA according to the behavior into the forwarding table for node PEA. This table is identified by the Mirror SID, which indicates node PEA's context. Using the forwarding entry for SIDA in this table, a packet with SIDA will be transmitted by PEB to the same destination as it is transmitted by PEA. For example, assume that the packet with SIDA is transmitted by PEA to CE2 through the forwarding behavior of the SIDA in PEA. The packet will be transmitted by PEB to the same CE2 through looking up the table identified by the Mirror SID.

After P1 as PLR receives the information <PEB, PEA, Mirror SID> and knows that PEB wants to protect SIDA of PEA, it computes an LFA for PEA assuming that PEA and PEB have a same anycast address. A Repair List RL is obtained based on the LFA. It is one of the followings:

- o RL = <Mirror SID> if the LFA is the next hop node to PEB along the shortest path to PEB; or
- o RL = <S1, ..., Sn, Mirror SID> if the LFA is a TI-LFA, where <S1, ..., Sn> is the TI-LFA Repair List to PEB computed by P1.

When PEA fails, P1 as PLR sends the packet with SIDA carried by the SR path to PEB, but encapsulates the packet before sending it by executing H.Encaps with the Repair List RL and a Source Address T.

Suppose that the packet received by P1 is represented by Pkt = (S, SIDA)Pkt0, where SA = S and DA = SIDA, and Pkt0 is the rest of the packet.

The execution of H.Encaps pushes an IPv6 header to Pkt and sets some fields in the outer and inner IPv6 header to produce an encapsulated packet Pkt'. Pkt' will be one of the followings:

- o Pkt' = (T, Mirror SID) (S, SIDA)Pkt0 if RL = <Mirror SID>; or
- o Pkt' = (T, S1)(Mirror SID, Sn, ..., S1; SL=n) (S, SIDA)Pkt0 if RL = <S1, ..., Sn, Mirror SID>.

When PEB receives the re-routed packet, which is (T, Mirror SID) (S, SIDA)Pkt0, it decapsulates the packet and forwards the decapsulated packet using the FIB table Tm identified by the Mirror SID as a variant of End.DT6 SID. The Mirror SID is called End.M.

It obtains the Mirror SID in the outer IPv6 header of the packet, removes this outer IPv6 header with all its extension headers, and

then processes the inner IPv6 packet (i.e., (S, SIDA)Pkt0, the packet without the outer IPv6 header). PEB finds the FIB table Tm for node PEA using the Mirror SID as the context ID, and submits the packet to this FIB table lookup and transmission to the same destination as PEA does.

The behavior of Mirror SID (End.M for short) is a variant of the End.DT6 behavior (refer to Section 4.6 of [RFC8986]). The End.M SID MUST be the last segment in an SR path, and a SID instance is associated with an IPv6 FIB table Tm.

When processing the Upper-Layer header of a packet matching a FIB entry locally instantiated as an End.M SID, N does the following:

```

S01. If (Upper-Layer header type == 41(IPv6) ) {
S02.   Remove the outer IPv6 header with all its extension headers
S03.   Set the packet's associated FIB table to Tm
S04.   Submit the packet to the egress IPv6 FIB lookup for
        transmission to the new destination
S05. } Else {
S06.   Process as per Section 4.1.1 of RFC8986
S07. }

```

3.2. Example

[Figure 2](#) shows an example of protecting egress PE3 of a SR path, which is from ingress PE1 to egress PE3.

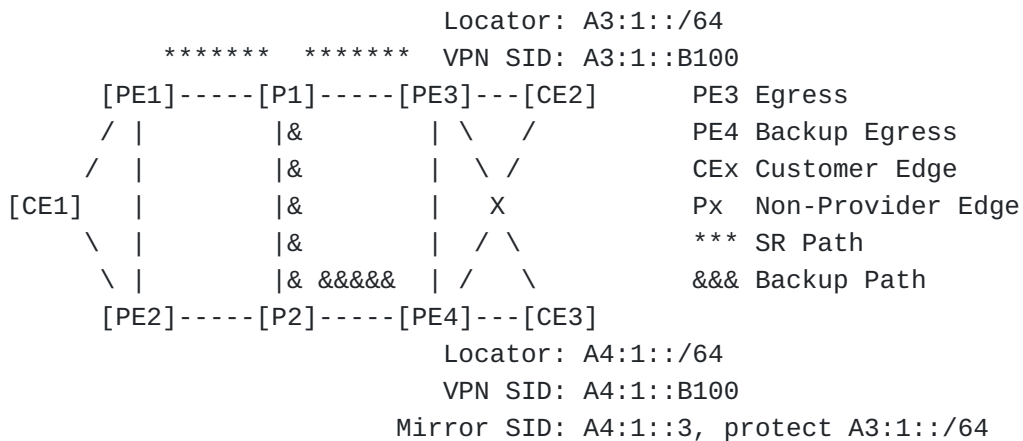


Figure 2: PE4 Protects Egress PE3 of SR Path

Where node P1's pre-computed backup path for PE3 is from P1 to PE4 via P2. In normal operations, after receiving a packet with destination PE3, P1 forwards the packet to PE3 according to its FIB. When PE3 receives the packet, it sends the packet to CE2.

When PE3 fails, P1 as PLR detects the failure through using a failure detection mechanism such as BFD and forwards the packet to PE4 via the backup path. When PE4 receives the packet, it sends the packet to the same CE2.

When P1's IGP converges on the failure of PE3, P1 as PLR needs to retent the route to PE3 for a period of time. Thus the backup path for PE3 will be used when the other nodes (such as PE1) still send the packet to PE3 via P1 since their IGP's do not converge on the failure.

In [Figure 2](#), Both CE2 and CE3 are dual home to PE3 and PE4. PE3 has a locator A3:1::/64 and a VPN SID A3:1::B100. PE4 has a locator A4:1::/64 and VPN SID A4:1::B100. A Mirror SID A4:1::3 is configured on PE4 for protecting PE3 with locator A3:1::/64.

After the configuration, PE4 advertises this information through an IGP LS (i.e., LSA in OSPF or LSP in IS-IS), which includes PE3's locator and Mirror SID A4:1::3. Every node in the SR domain will receive this IGP LS, which indicates that PE4 wants to protect PE3 (indicated by PE3's locator) with Mirror SID A4:1::3.

When PE4 (e.g., BGP on PE4) receives a prefix whose VPN SID belongs to PE3 that is protected by PE4 through Mirror SID A4:1::3, it finds PE4's VPN SID corresponding to PE3's VPN SID. For example, local PE4 has Prefix 1.1.1.1 with VPN SID A4:1::B100, when PE4 receives prefix 1.1.1.1 with remote PE3's VPN SID A3:1::B100, it knows that they are for the same VPN.

The forwarding behaviors for these two VPN SIDs are the same from function's point of view. If the behavior for PE3's VPN SID in PE3 forwards the packet with it to CE2, then the behavior for PE4's VPN SID in PE4 forwards the packet to the same CE2; and vice versa. PE4 creates a forwarding entry for PE3's VPN SID A3:1::B100 in the FIB table identified by Mirror SID A4:1::3 according to the forwarding behavior for PE4's VPN SID A4:1::B100.

Node P1's pre-computed backup path for destination PE3 is from P1 to PE4 having mirror SID A4:1::3. When P1 receives a packet destined to PE3's VPN SID A3:1::B100, in normal operations, it forwards the packet with source A1:1:: and destination PE3's VPN SID A3:1::B100 according to the FIB using the destination PE3's VPN SID A3:1::B100.

When PE3 fails, P1 as PLR sends the packet to PE4 via the backup path pre-computed. P1 encapsulates the packet using H.Encaps before sending it to PE4.

Suppose that the packet received by P1 is represented by $\text{Pkt} = (\text{SA} = \text{A1:1::}, \text{DA} = \text{A3:1::B100})\text{Pkt}_0$, where $\text{DA} = \text{A3:1::B100}$ is PE3's VPN

SID, and Pkt0 is the rest of the packet. The encapsulated packet Pkt' will be one of the followings:

- o Pkt' = (T, Mirror SID A4:1::3) (A1:1::, A3:1::B100)Pkt0 if backup path not via PE3; or (otherwise)
- o Pkt' = (T, S1)(Mirror SID A4:1::3, Sn, ..., S1; SL=n) (A1:1::, A3:1::B100)Pkt0.

where T is a Source Address, <S1, ..., Sn> is the TI-LFA Repair List to PE4 computed by P1 when the backup path to PE4 goes through PE3.

When PE4 receives the re-routed packet, it decapsulates the packet and forwards the decapsulated packet by executing End.DT6 behavior for an End.DT6 SID instance. The SID instance is End.M, the Mirror SID that is associated with the IPv6 FIB table for PE3. The packet received by PE4 is (T, Mirror SID A4:1::3) (A1:1::, PE3's VPN SID A3:1::B100)Pkt0.

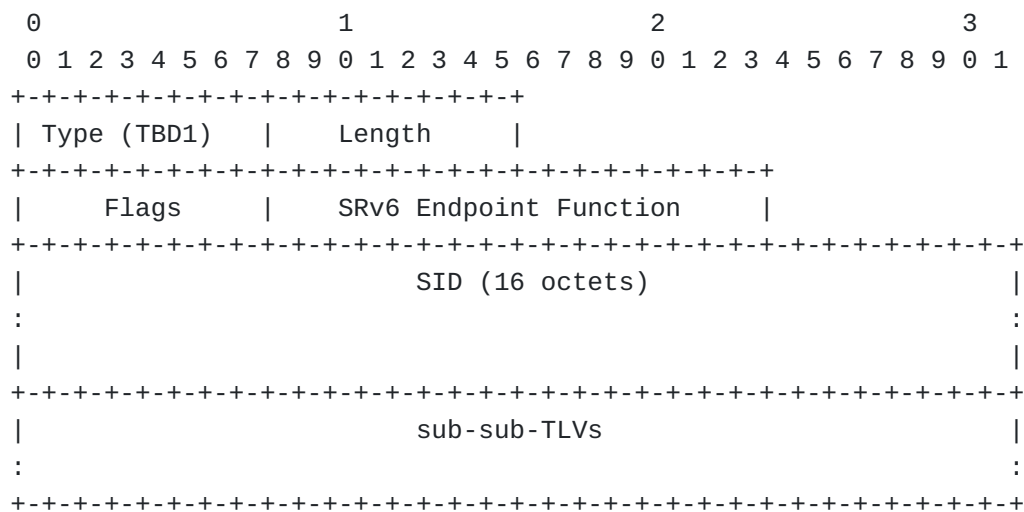
PE4 obtains Mirror SID A4:1::3 in the outer IPv6 header of the packet, removes this outer IPv6 header, and then processes the inner IPv6 packet (A1:1::, A3:1::B100)Pkt0. It finds the FIB table for PE3 using Mirror SID A4:1::3 as the context ID, gets the forwarding entry for PE3's VPN SID A3:1::B100 from the table, and forwards the packet to CE2 using the entry.

4. Extensions to IGP for Egress Protection

This section describes extensions to IS-IS and OSPF for advertising the information about SRv6 path egress protection.

4.1. Extensions to IS-IS

A new sub-TLV, called IS-IS SRv6 Mirror SID sub-TLV, is defined. It is used in the SRv6 Locator TLV defined in [[RFC9352](#)] to advertise SRv6 Mirror SID and the locators of the nodes to be protected. The SRv6 Mirror SID inherit the topology/algorithm from the parent locator. The format of the sub-TLV is illustrated below.



Type: TBD1 (suggested value 8) is to be assigned by IANA.

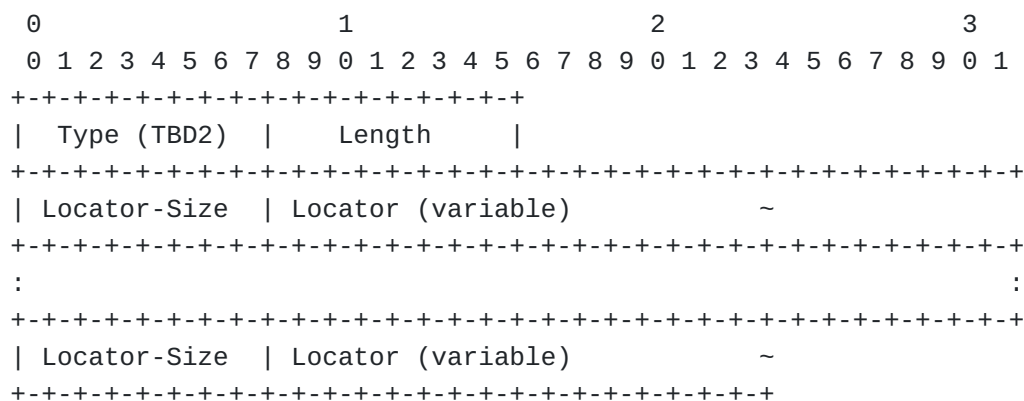
Length: variable.

Flags: 1 octet. No flags are currently defined.

SRv6 Endpoint Function: 2 octets. It contains the endpoint function 74 for Mirror SID.

SID: 16 octets. This field contains the SRv6 Mirror SID to be advertised.

A protected locators sub-sub-TLV is defined and used to carry the Locators of the egress nodes to be protected by the SRv6 mirror SID. It has the following format.



Type: TBD2 (suggested value 1) is to be assigned by IANA.

Locator-Size: 1 octet. Number of bits (1 - 128) in the Locator field.

When node B advertises that B wants to protect node A with a Mirror SID through an LSP, the LSP contains an IS-IS SRv6 Mirror SID sub-TLV, which includes the Mirror SID and node A's locator in an IS-IS Protected locators sub-sub-TLV.

Similarly, a new sub-TLV, called OSPF Mirror SID sub-TLV, is defined. It is used to advertise SRv6 Mirror SID and the locators of the nodes to be protected. Its format is illustrated below.



Type:

TBD4 (suggested value 8) is to be assigned by IANA.

Length: variable.

Flags: 1 octet. No flags are currently defined.

Reserved: 1 octet. It MUST be set to zero for transmission and ignored on reception.

SRv6 Endpoint Function: 2 octets. It contains the endpoint function 74 for End.M SID.

SID: 16 octets. This field contains the SRv6 Mirror SID to be advertised.

A protected locators sub-TLV is defined and used to carry the locators of the nodes to be protected by the SRv6 Mirror SID. It has the following format.

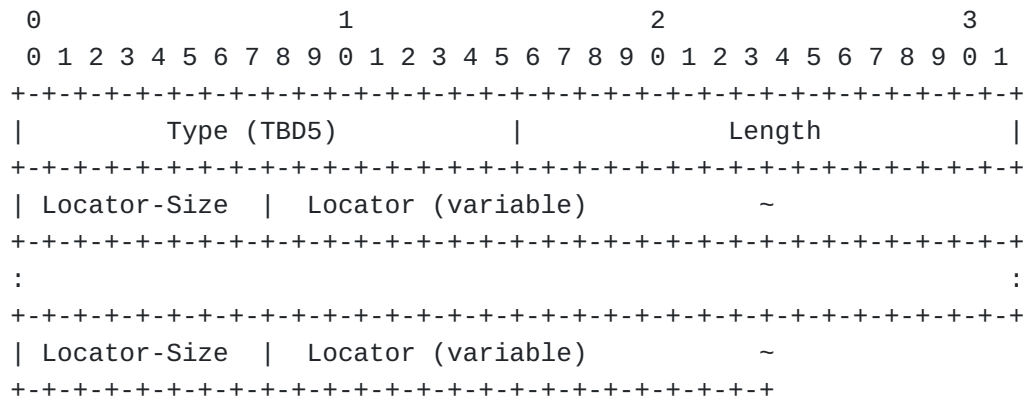


Figure 6: OSPF Protected Locators sub-TLV

Type: TBD5 (suggested value 1) is to be assigned by IANA.

Length: variable.

Locator-Size: 1 octet. Number of bits (1 - 128) in the Locator field.

Locator: 1-16 octets. This field encodes an SRv6 Locator of an egress node to be protected by the SRv6 mirror SID. The Locator is encoded in the minimal number of octets for the given number of bits.

5. Security Considerations

The security about the egress protection is described in details in [RFC8679]. The extensions to OSPF and IS-IS described in this document for SRv6 path egress protection should not cause extra security issues.

6. IANA Considerations

6.1. SRv6 Endpoint Behaviors

Under sub-registry "SRv6 Endpoint Behaviors" [RFC8986], IANA has assigned the following for End.M Endpoint Behavior:

Value	Hex	Endpoint behavior	Reference
74	0x004A	End.M (Mirror SID)	This document

6.2. IS-IS

Under "IS-IS Sub-TLVs for TLVs Advertising Prefix Reachability registry", IANA is requested to add the following new Sub-TLV:

Type	Description	Reference
8	SRv6 Mirror SID	This document

IANA is requested to create and maintain a new registry for sub-sub-TLVs of the SRv6 Mirror SID Sub-TLV. The suggested registry name is

o Sub-Sub-TLVs for SRv6 Mirror SID Sub-TLV

Initial values for the registry are given below. The future assignments are to be made through IETF Review [RFC5226].

Value	Sub-Sub-TLV Name	Definition
0	Reserved	
1	Protected Locators Sub-Sub-TLV	This Document
2-255	Unassigned	

6.3. OSPFv3

Under registry "OSPFv3 Locator LSA Sub-TLVs" [I-D.ietf-lsr-ospfv3-srv6-extensions], IANA is requested to assign the following new Sub-TLVs:

Sub-TLV Type	Sub-TLV Name	Reference
8	SRv6 Mirror SID Sub-TLV	This document
11	Protected Locators Sub-TLV	This document

7. References

7.1. Normative References

- [I-D.ietf-lsr-ospfv3-srv6-extensions] Li, Z., Hu, Z., Talaulikar, K., and P. Psenak, "OSPFv3 Extensions for SRv6", Work in Progress, Internet-Draft, draft-ietf-lsr-ospfv3-srv6-extensions-09, 14 January 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsr-ospfv3-srv6-extensions-09>>.
- [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>>.
- [RFC8400] Chen, H., Liu, A., Saad, T., Xu, F., and L. Huang, "Extensions to RSVP-TE for Label Switched Path (LSP) Egress Protection", RFC 8400, DOI 10.17487/RFC8400, June 2018, <<https://www.rfc-editor.org/info/rfc8400>>.
- [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>>.
- [RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", RFC 8667, DOI 10.17487/RFC8667, December 2019, <<https://www.rfc-editor.org/info/rfc8667>>.
- [RFC8679] Shen, Y., Jeganathan, M., Decraene, B., Gredler, H., Michel, C., and H. Chen, "MPLS Egress Protection Framework", RFC 8679, DOI 10.17487/RFC8679, December 2019, <<https://www.rfc-editor.org/info/rfc8679>>.
- [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>>.

[RFC9352]

Psenak, P., Ed., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Segment Routing over the IPv6 Data Plane", RFC 9352, DOI 10.17487/RFC9352, February 2023, <<https://www.rfc-editor.org/info/rfc9352>>.

7.2. Informative References

[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-09, 23 December 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-rtgwg-segment-routing-ti-lfa-09>>.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, DOI 10.17487/RFC5226, May 2008, <<https://www.rfc-editor.org/info/rfc5226>>.

Acknowledgments

The authors would like to thank Acee Lindem, Peter Psenak, Yimin Shen, Zhenqiang Li, Alexander Vainshtein, Greg Mirsky, Bruno Decraene, Jeff Tantsura, Chris Bowers and Ketan Talaulikar for their comments to this work.

Authors' Addresses

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

Email: huzhibo@huawei.com

Huaimo Chen
Futurewei
Boston, MA,
United States of America

Email: Huaimo.chen@futurewei.com

Huanan Chen
China Telecom
109, West Zhongshan Road, Tianhe District

Guangzhou
510000
China

Email: chenhn8.gd@chinatelecom.cn

Peng Wu
Huawei
Huawei Bld., No.156 Beiqing Rd.
Beijing
100095
China

Email: baggio.wupeng@huawei.com

Mehmet Toy
Verizon
United States of America

Email: mehmet.toy@verizon.com

Chang Cao
China Unicom

Email: caoc15@chinaunicom.cn

Tao He
China Unicom

Email: het21@chinaunicom.cn

Lei Liu
Fujitsu
United States of America

Email: liulei.kddi@gmail.com

Xufeng Liu
IBM Corporation
United States of America

Email: xufeng.liu.ietf@gmail.com