Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 11 October 2026

T. He
China Unicom
Z. Hu
Huawei
H. Chen
Futurewei
M. Toy
Verizon
C. Cao
China Unicom

9 April 2026

SRv6 Path Egress Protection
draft-ietf-rtgwg-srv6-egress-protection-23

Abstract

TI-LFA specifies fast protections for transit nodes and links of an SR path. However, it does not present any fast protections for the egress node of the SR path. This document describes protocol extensions for fast protecting the egress node and link of a Segment Routing for IPv6 (SRv6) path. The solution uses IGP extensions and a Mirror SID (End.M) behavior to steer traffic to a protector egress upon failure of the primary egress.

This document operates within a single link-state IGP area/level and uses IS-IS/OSPFv3 to advertise a Mirror SID and the protected locators for egress node/link protection. While the mechanism can protect traffic whose active segment at the egress is a Service SID (e.g., VPN SID), it is not suitable for large-scale deployments with a high cardinality of VPN/service instances or random multi-homing patterns, because the amount of egress-protection information to be flooded in the IGP increases and may impact convergence and control-plane load.

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.

This Internet-Draft will expire on 11 October 2026.

Copyright Notice

Table of Contents

1. Introduction 3
2. Terminologies 4
3. SRv6 Path Egress Protection 5
  3.1. Mechanism 5
    3.1.1. Egress Node Protection 6
    3.1.2. Egress Link Protection 9
  3.2. Example 9
  3.3. Operational Guidelines 12
4. Extensions to IGP for Egress Protection 12
  4.1. Extensions to IS-IS 12
  4.2. Extensions to OSPFv3 14
5. Security Considerations 16
6. IANA Considerations 19
  6.1. SRv6 Endpoint Behaviors 19
  6.2. IS-IS 19
  6.3. OSPFv3 20
7. References 20
  7.1. Normative References 20
  7.2. Informative References 22
Acknowledgments 22

Contributors' Addresses 22
Authors' Addresses 23

1. Introduction

[RFC9855] specifies fast protections for nodes and links that are within a link-state IGP area. In other words, it specifies fast protections for transit nodes and links of an SR path, but does not describe any fast protections for the egress node or link of an SR path. While TI-LFA provides fast protection for transit nodes and links within an IGP area, it does not address egress node protection because the egress node is the endpoint of the SR path. TI-LFA relies on pre-computed backup paths that bypass failed nodes or links while maintaining the same destination. However, when the egress node itself fails, the destination becomes unreachable through the primary path. This document addresses this gap by introducing a Mirror SID mechanism that allows a backup egress node to take over the forwarding behavior of the failed egress node, effectively extending protection to the network edge.

[RFC8400] and [RFC8679] specify fast protections for egress node(s) and link(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 any fast protection for the egress node and link of a Segment Routing for IPv6 (SRv6) path or tunnel.

For an SRv6 path from an ingress node to an egress node, the fast protection for the egress node and link of the path can be achieved through using 1 + 1 global protection. This solution uses more network resources and makes operation complex. A backup SRv6 path from the ingress node to a backup egress node is set up. A CE is dual-homed to the egress node and the backup egress node. A SID of the egress node is used to forward the traffic to the CE. This same SID is configured on the backup egress node to forward the traffic to the same CE. Both paths transmit the traffic to the same CE, which selects one. The CE selects the traffic from the egress node if the egress node and link work well; otherwise (i.e., the egress node or link failed), the CE selects the traffic from the backup egress node.

This document presents a solution which provides fast protections for the egress node and link of an SRv6 path through extending IGP and using Mirror SID. Compared to 1 + 1 global protection, this solution is more efficient and the operation on it is simpler.

The solution is scoped to a single link-state IGP area/level. It relies on IGP to distribute the tuple <PEB, PEA, Mirror SID> with the protected locators. The forwarding behavior for Service SIDs anchored on PEA may be obtained by the protector via existing means

(e.g., [RFC9252]) or configuration, but this document does not introduce any per-service signaling in the IGP. Furthermore, the approach is applicable to modest numbers of protected services; large-scale deployments with many VPN/service instances or random multi-homing are not recommended due to IGP scaling considerations.

2. Terminologies

The following terminologies are used in this document.

BFD: Bidirectional Forwarding Detection

BGP: Border Gateway Protocol

CE: Customer Edge

DA: Destination Address

Egress link: A link from an egress node to another domain [RFC8679]

Egress node: A domain exit node on an SRv6 path

FIB: Forwarding Information Base

IGP: Interior Gateway Protocol

IS-IS: Intermediate System to Intermediate System

L3VPN: Layer 3 VPN

LFA: Loop-Free Alternate

LS: Link State, which is LSA in OSPF/OSPFv3 or LSP in IS-IS

LSA: Link State Advertisement in OSPF/OSPFv3

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

OSPF: Open Shortest Path First

OSPFv3: Open Shortest Path First version 3

P2MP: Point-to-MultiPoint

P2P: Point-to-Point

PDU: Protocol Data Unit

PE: Provider Edge

PLR: Point of Local Repair

RL: Repair List

SA: Source Address

SID: Segment Identifier

SR: Segment Routing

SR path: An SR path in this document is the active path of an SR Policy [RFC9256]

SRv6: SR for IPv6

SRv6 path: An SRv6 path in this document is the active path of an SR Policy with SRv6 SIDs [RFC9256]

TE: Traffic Engineering

TI-LFA: Topology Independent LFA

VPN: Virtual Private Network

3. SRv6 Path Egress Protection

This section describes the mechanism of SRv6 path egress protection and illustrates it through an example.

All advertisements and computations in this section are confined to a single link-state IGP area/level for SRv6.

3.1. Mechanism

Figure 1 is used to explain the mechanism of SRv6 path egress node and egress link protection.

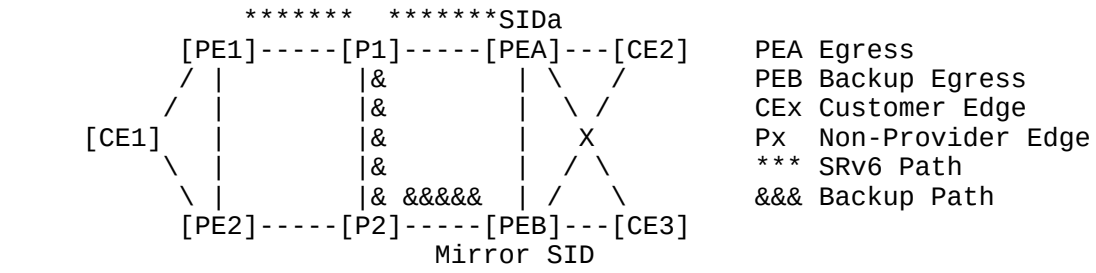

Figure 1: PEB Protects Egress PEA of SRv6 Path

3.1.1. Egress Node Protection

Desired Pathways in Figure 1:

Node PEA in Figure 1 is the egress node (aka egress) of the SRv6 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 node (aka protector or backup egress) to provide the fast protection for the egress node (aka primary egress node) PEA. Node P1 is the direct previous/upstream endpoint of egress node PEA and acts as PLR (refer to [RFC9855]) to support the fast protection for PEA.

Steps in Creating the Pathways:

Step 1: Normal Pathway Set-up

Normal path set-up establishes the SR path from ingress PE1 to egress PEA via P1. Ingress PE1 imports the traffic from CE1 into the SR path and egress PEA delivers the traffic from the SRv6 path to CE2.

Step 2: Backup Pathway Set-up

Step 2a: PEB Announces to Protect PEA

When PEB is selected as a backup egress node to protect the egress node PEA, a SRv6 Mirror SID (refer to Section 5.1 of [RFC8402]) is configured on PEB to protect PEA. PEB MUST advertise this information through IGP, which includes the Mirror SID and the egress PEA. The information is represented by <PEB, PEA, Mirror SID>, together with the protected locators. This IGP signaling does not enumerate per-service entries.

Step 2b: PEB Gets Forwarding Behavior of PEA

After PEA receives the information <PEB, PEA, Mirror SID>, it may provide to PEB the forwarding behavior for the active SRv6 segment SIDa at PEA by existing means. When SIDa is a Service SID (e.g., a VPN SID) anchored on PEA, PEB may learn its forwarding behavior via the BGP-based overlay as per [RFC9252] or by configuration; this document does not introduce per-service signaling in IGP. This enables PEB to reproduce the egress behavior for packets whose active segment at PEA is a Service SID, without requiring IGP to flood per-service state.

Step 2c: PEB Creates FIB for PEA

When PEB gets the forwarding behavior of SIDa of PEA, it MUST add a forwarding entry for SIDa into the forwarding table identified by the Mirror SID (the PEA context). This supports Service SID semantics at the protector. However, for large numbers of Service SIDs, operators SHOULD avoid deployments where protection requires fine-grained per-service modeling in IGP, as it may increase IGP flooding and affect convergence.

Step 2d: P1 as PLR Prepares to Protect PEA by PEB

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 the same anycast address. A Repair List RL (or say backup path) is obtained based on the LFA. It is one of the following:

- RL = <Mirror SID> if the LFA is the next hop node to PEB along the shortest path to PEB; or

- 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.

Step 3: Backup Path Is Engaged upon PEA Failure

Step 3a: P1 Detects PEA Failure via BFD or Other Mechanisms

Step 3b: P1 Sends Packet with SIDa to Backup Egress PEB

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

P1 as PLR needs to retain the route to PEA for a period of time after its IGP converges on the failure of PEA. Thus the backup path for PEA will be used when the other nodes (such as PE1) still send the packet to PEA via P1 since their IGPs do not converge on the failure.

Suppose that the packet received by P1 is represented by Pkt = (S, SID-P1)(SIDa,SID-P1; SL=1)Pkt0, where SA = S and DA = SID-P1 (i.e., SID of P1), and Pkt0 is the rest of the packet. P1 sets DA to SIDa, updates SL and executes H.Encaps.

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 following:

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

Step 3c: PEB Decapsulates Packet and Forwards It

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.

When a node processes a packet with an End.M SID as the destination, it MUST perform the following steps in order: 1. Verify that the End.M SID is locally instantiated. If not, process according to Section 4.1.1 of [RFC8986]. 2. Remove the outer IPv6 header with all its extension headers. 3. Identify the FIB table associated with the End.M SID context. 4. Submit the inner packet to the identified FIB table for forwarding. If any of these steps fail, the packet MUST be dropped and an appropriate error counter SHOULD be incremented.

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.

3.1.2. Egress Link Protection

Egress link protection is similar to egress node protection for SRv6 [RFC8679]. When the egress link from egress node PEA to CE2 fails, PEA acting as a PLR reroutes the traffic to backup egress node PEB via a backup path. Specifically, PEA as a PLR pre-computes a Repair List RL (or say backup path) toward PEB after receiving <PEB, PEA, Mirror SID> and knowing that PEB wants to protect SIDa of PEA. When the link fails, PEA as PLR sends the packet with SIDa by executing H.Encaps with the Repair List RL. All operations occur within the same IGP flooding scope; no per-service signaling is introduced in IGP.

3.2. Example

Figure 2 shows an example of fast protecting egress node PE3 of an SRv6 path, which is from ingress node PE1 to egress node PE3.

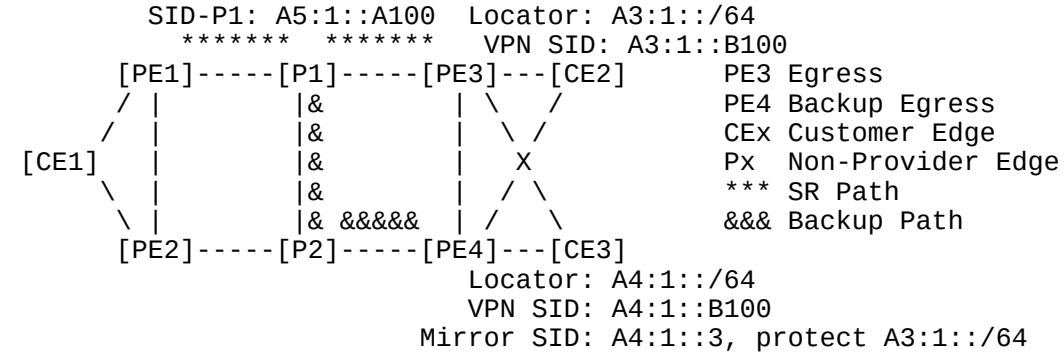

Figure 2: PE4 Protects Egress PE3 of SR Path

Desired Pathways in Figure 2:

Node P1’s pre-computed backup path for PE3 is from P1 to PE4 via P2. In normal operations, after receiving a packet with SRv6 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 retain 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 IGPs do not converge on the failure.

In Figure 2, Both CE2 and CE3 are dual-homed 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. P1 has SID-P1 = A5:1::A100.

Steps in Creating the Pathways:

Step 1: Normal Pathway Set-up [PEB is PE4, PEA is PE3]

Step 2: Backup Pathway Set-up

Step 2a: PE4 (aka PEB) Announces to Protect PE3 (aka PEA)

After the configuration, PE4 advertises this information through an IGP LS (i.e., LSA in OSPFv3 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.

Step 2b: PE4 (aka PEB) Gets Forwarding Behavior of PE3 (aka PEA)

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.

Step 2c: PE4 (aka PEB) Creates FIB for PE3 (aka PEA)

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.

Step 2d: P1 Prepares to Protect PE3 (aka PEA) by PE4 (aka PEB)

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.

Step 3: Backup Path Is Engaged upon PE3 (aka PEA) Failure

Step 3a: P1 Detects PE3 (aka PEA) Failure via BFD

Step 3b: P1 Sends Packet with SIDa to Backup Egress PE4 (aka PEB)

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 Pkt = (SA=A1:1::,DA=A5:1::A100)(SIDa=A3:1::B100,SID-P1=A5:1::A100;SL=1) Pkt0, where DA = A5:1::A100 is P1’s SID, A3:1::B100 is PE3’s VPN SID, and Pkt0 is the rest of the packet. P1 sets DA to A3:1::B100, updates SL, and encapsulates the packet. The encapsulated packet Pkt’ will be one of the following:

- Pkt’ = (T, Mirror SID A4:1::3) (A1:1::, A3:1::B100)Pkt0 if the LFA is the next hop node to PE4 along the shortest path to PE4; or (otherwise)

- 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.

Step 3c: PE4 (aka PEB) Decapsulates Packet and Forwards It

When PE4 receives the re-routed packet, it decapsulates the packet and forwards the decapsulated packet by executing the behavior of End.M for 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.

Note:This example demonstrates that a Service SID (e.g., a VPN SID) can be preserved at the protector via the Mirror-SID context. However, at large scale (many VPN/service instances and/or random multi-homing of services across multiple protectors), the amount of egress-protection information to be flooded in IGP increases and may affect convergence; such deployments are not recommended for this IGP-based mechanism. Operators SHOULD consolidate protectors per egress and limit per-service granularity in IGP.

3.3. Operational Guidelines

Protector Consolidation: Prefer a single protector PEB per PEA within the IGP area/level to minimize the number of SRv6 <PEB, PEA, Mirror SID> advertisements.

Limit Granularity in IGP: Do not attempt to enumerate per-service entries in IGP; use locator-level protection only.

Service Behavior Acquisition: Learning Service-SID behaviors at the protector (e.g., via BGP per [RFC9352] or configuration) is an implementation choice and does not alter the IGP flooding scope.

Applicability Thresholds: When the protected service count per egress or the number of protection relationships grows large-especially with random multi-homing-the IGP control-plane load and convergence may be adversely affected; such deployments are not recommended for this mechanism.

4. Extensions to IGP for Egress Protection

This section describes extensions to IS-IS and OSPFv3 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 inherits the topology/algorithm from the parent locator. The format of the sub-TLV is illustrated below.

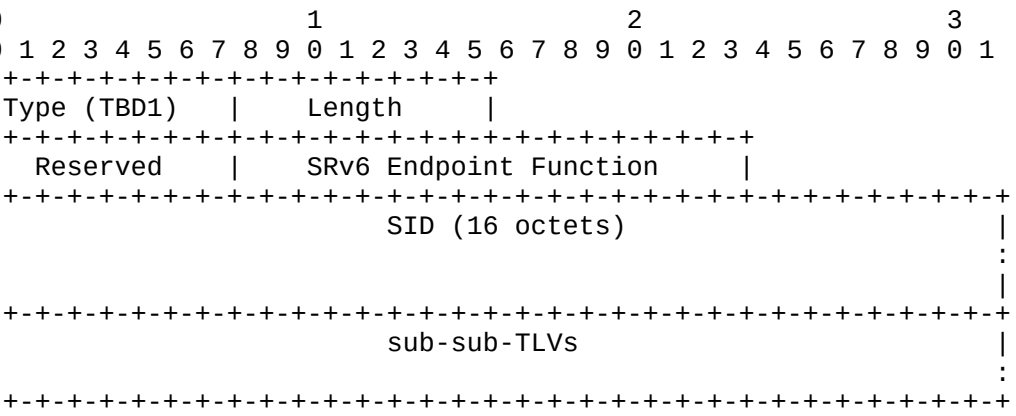

Figure 3: IS-IS SRv6 Mirror SID sub-TLV

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

Length: 1 octet. Its value MUST NOT be less than 23. 23 is 19 (i.e., the size of Reserved, SRv6 Endpoint Function and SID) plus 4 (i.e., the minimum size of a IS-IS protected locators sub-sub-TLV). The entire IS-IS SRv6 Mirror SID sub-TLV MUST be ignored if the length is less than 23.

Reserved: 1 octet. This octet MUST be set to zero on transmit, and ignored on receipt.

SRv6 Endpoint Function: 2 octets. It MUST contain the endpoint function 74 for Mirror SID. The entire IS-IS SRv6 Mirror SID sub-TLV MUST be ignored if it does not contain the endpoint function 74.

SID: 16 octets. This field contains the SRv6 Mirror SID to be advertised. It MUST NOT be zero (0). The entire IS-IS SRv6 Mirror SID sub-TLV MUST be ignored if it contains zero (0).

A protected locators sub-sub-TLV is defined and used to carry the Locators of the egress node to be protected by the SRv6 mirror SID. The IS-IS SRv6 Mirror SID sub-TLV MUST include one IS-IS protected locators sub-sub-TLV. It has the following format.

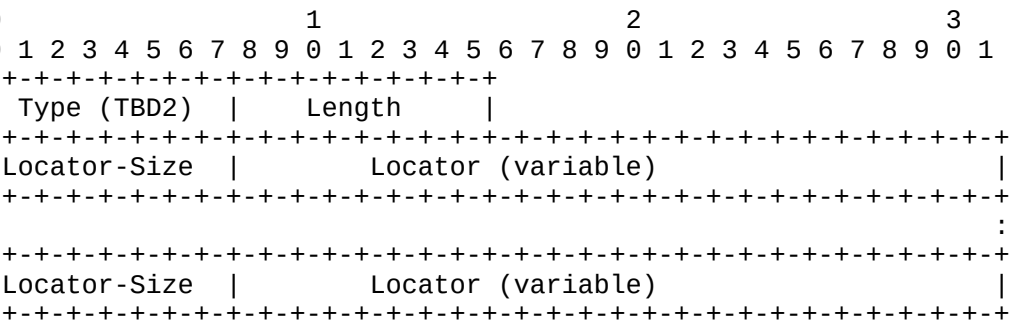

Figure 4: IS-IS Protected Locators sub-sub-TLV

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

Length: 1 octet. Its value MUST NOT be less than 2. The entire IS-IS SRv6 Mirror SID sub-TLV MUST be ignored if the length is less than 2.

Locator-Size: 1 octet. Number of bits in the Locator field, which MUST be from the range (1-128). The entire IS-IS SRv6 Mirror SID sub-TLV MUST be ignored if the Locator-Size is outside this range.

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. Trailing bits MUST be set to zero and ignored when received.

When node B advertises that B wants to protect node A with a Mirror SID through an LSP, the LSP MUST have an SRv6 Locator TLV containing an IS-IS SRv6 Mirror SID sub-TLV, which includes the Mirror SID and node A’s locators in an IS-IS Protected locators sub-sub-TLV.

Note:The IS-IS SRv6 Mirror SID sub-TLV MUST include exactly one "Protected Locators" sub-sub-TLV and MUST NOT carry per-service (e.g., VPN/Service-SID) enumerations. This document does not define any IGP encoding to list individual services; attempting to do so at large scale is not suitable due to IGP flooding and convergence considerations.

4.2. Extensions to OSPFv3

Similarly, a new sub-TLV, called OSPFv3 Mirror SID sub-TLV, is defined. It is used in the SRv6 Locator TLV defined in [RFC9513] to advertise SRv6 Mirror SID and the locators of the nodes to be protected. Its format is illustrated below.

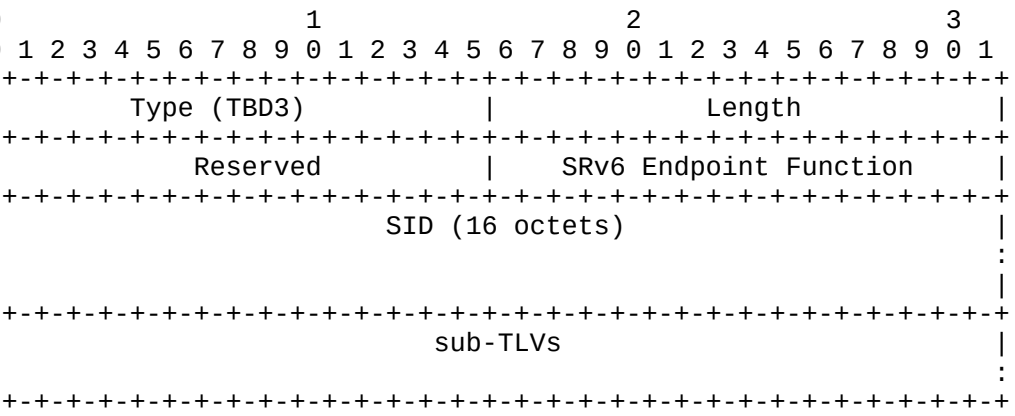

Figure 5: OSPFv3 SRv6 Mirror SID sub-TLV

Type: TBD3 (suggested value 8) is to be assigned by IANA.

Length: 2 octets. Its value MUST NOT be less than 26. 26 is 20 (i.e., the size of Reserved, SRv6 Endpoint Function and SID) plus 6 (i.e., the minimum size of a OSPFv3 protected locators sub-TLV). The entire OSPFv3 SRv6 Mirror SID sub-TLV MUST be ignored if the length is less than 26.

Reserved: 2 octets. It MUST be set to zero for transmission and ignored on reception.

SRv6 Endpoint Function: 2 octets. It MUST contain the endpoint function 74 for End.M SID. The entire OSPFv3 SRv6 Mirror SID sub-TLV MUST be ignored if it does not contain the endpoint function 74.

SID: 16 octets. This field contains the SRv6 Mirror SID to be advertised. It MUST NOT be zero (0). The entire OSPFv3 SRv6 Mirror SID sub-TLV MUST be ignored if it contains zero (0).

A protected locators sub-TLV is defined and used to carry the locators of the node to be protected by the SRv6 Mirror SID. The OSPFv3 SRv6 Mirror SID sub-TLV MUST include one OSPFv3 protected locators sub-TLV. It has the following format.

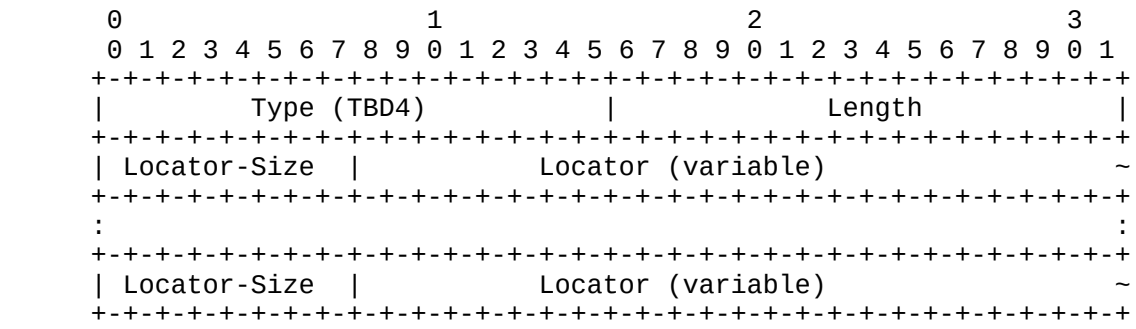

Figure 6: OSPFv3 Protected Locators sub-TLV

Type: TBD4 (suggested value 1) is to be assigned by IANA.

Length: 2 octets. Its value MUST NOT be less than 2. The entire OSPFv3 SRv6 Mirror SID sub-TLV MUST be ignored if the Length is less than 2.

Locator-Size: 1 octet. Number of bits (1 - 128) in the Locator field. Number of bits in the Locator field, which MUST be from the range (1-128). The entire OSPFv3 SRv6 Mirror SID sub-TLV MUST be ignored if the Locator-Size is outside this range.

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. Trailing bits MUST be set to zero and ignored when received.

When node B advertises that B wants to protect node A with a Mirror SID through an LSA, the LSA MUST have an SRv6 Locator TLV containing an OSPFv3 SRv6 Mirror SID sub-TLV, which includes the Mirror SID and node A’s locators in an OSPFv3 Protected Locators sub-TLV.

Note:The OSPFv3 SRv6 Mirror SID sub-TLV MUST include exactly one "Protected Locators" sub-TLV and MUST NOT carry per-service (e.g., VPN/Service-SID) enumerations for the same reasons as above.

5. Security Considerations

The egress protection specified in this document involves rerouting traffic around an egress node or link failure, via a backup path from a PLR to a backup egress node. The forwarding performed by the nodes in the data plane is anticipated, as part of the planning of egress protection.

The extensions to control plane protocol IS-IS or OSPFv3 are used to support the egress protection on the nodes in an OSPFv3 or IS-IS area. The area is in a single administrative domain.

In addition, the PLR and backup egress node are located close to the egress node, which is in the same administrative domain.

Security concerns for IS-IS are addressed in [ISO10589], [RFC5304] and [RFC5310]. While IS-IS is deployed under a single administrative domain, there can be deployments where potential attackers have access to one or more networks in the IS-IS routing domain. In these deployments, the stronger authentication mechanisms defined in the aforementioned documents SHOULD be used.

Security concerns for OSPFv3 are described in [RFC5340] and [RFC8362]. While OSPFv3 is under a single administrative domain, there can be deployments where potential attackers have access to one or more networks in the OSPFv3 routing domain. In these deployments, stronger authentication mechanisms such as those specified in [RFC4552] and [RFC7166] SHOULD be used.

SRv6-Specific Security Considerations:The Mirror SID mechanism introduces the following SRv6-specific security considerations:

- Mirror SID Authentication: Since the Mirror SID is advertised through IGP, it is essential that IGP advertisements are authenticated to prevent malicious nodes from advertising counterfeit Mirror SIDs. Implementations SHOULD support the cryptographic authentication mechanisms specified in [RFC5304] for IS-IS and [RFC4552] for OSPFv3.

- Context Isolation: The FIB table identified by the Mirror SID MUST be properly isolated to prevent traffic from one protected egress node from being forwarded using the context of another egress node. Implementations MUST ensure that the context identified by a Mirror SID is only used for packets explicitly addressed to that Mirror SID.

Security attacks may sometimes come from a customer domain. Such attacks are not introduced by the egress protection in this document and may occur regardless of the existence of egress protection. In one possible case, the egress link between an egress node and a CE could become a point of attack. An attacker that gains control of the CE might use it to simulate link failures and trigger constant and cascading activities in the network. If egress link protection is in place, egress link protection activities may also be triggered. As a general solution to defeat the attack, a damping mechanism SHOULD be used by the egress node to promptly suppress the services associated with the link or CE. The egress node would stop

delivering the services to CE, essentially detaching them from the network and eliminating the effect of the simulated link failures. All protocol extensions operate within a single link-state IGP area/ level; no per-service signaling is introduced in IGP, and references to BGP concern only how a protector may learn forwarding behavior.When a protector learns per-service forwarding behavior via mechanisms outside the IGP (e.g., BGP as per [RFC9252] or local configuration), it SHOULD validate that the behavior is authorized and consistent with the protected egress node’s capabilities advertised through those same external mechanisms. This document does not introduce per-service signaling in the IGP for this validation.

The following threats are addressed by corresponding mechanisms:

-- Unauthorized Rerouting: An attacker could attempt to trigger unnecessary rerouting by advertising false failure information. This is mitigated by:

* Using BFD or other reliable failure detection mechanisms.

* Authenticating IGP advertisements.

* Implementing damping mechanisms to prevent flapping.

-- Traffic Interception: An attacker could attempt to become a protector to intercept traffic. This is mitigated by:

* Limiting protector selection to trusted nodes within the same administrative domain.

* Authenticating IGP advertisements of Mirror SIDs.

* Monitoring for unexpected protector advertisements.

-- Denial of Service: An attacker could attempt to overwhelm the protector with traffic. This is mitigated by:

* Implementing rate limiting at the protector.

* Using proper FIB table isolation.

* Monitoring traffic patterns.

-- Control Plane Overload: An attacker could attempt to flood the IGP with excessive protection advertisements, causing control plane overload and convergence issues.This is mitigated by:

* Limit the number of protection relationships per node.

* Implement IGP flooding rate limiting.

* Use the operational guidelines in Section 3.3 to limit the scale of deployments.

* Monitor IGP database size and convergence times.

6. IANA Considerations

6.1. SRv6 Endpoint Behaviors

Under 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
TBD	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 suggested 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 SRv6 Locator LSA Sub-TLVs" [RFC9513], IANA is requested to assign the following new Sub-TLVs:

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

7. References

7.1. Normative References

[ISO10589] ISO, "Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002, November 2002.

[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>.

[RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, DOI 10.17487/RFC4552, June 2006, <https://www.rfc-editor.org/info/rfc4552>.

[RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <https://www.rfc-editor.org/info/rfc5304>.

[RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <https://www.rfc-editor.org/info/rfc5310>.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <https://www.rfc-editor.org/info/rfc5340>.

[RFC7166] Bhatia, M., Manral, V., and A. Lindem, "Supporting Authentication Trailer for OSPFv3", RFC 7166, DOI 10.17487/RFC7166, March 2014, <https://www.rfc-editor.org/info/rfc7166>.

[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>.

[RFC8362] Lindem, A., Roy, A., Goethals, D., Reddy Vallem, V., and F. Baker, "OSPFv3 Link State Advertisement (LSA) Extensibility", RFC 8362, DOI 10.17487/RFC8362, April 2018, <https://www.rfc-editor.org/info/rfc8362>.

[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>.

[RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <https://www.rfc-editor.org/info/rfc9256>.

[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>.

[RFC9513] Li, Z., Hu, Z., Talaulikar, K., Ed., and P. Psenak, "OSPFv3 Extensions for Segment Routing over IPv6 (SRv6)", RFC 9513, DOI 10.17487/RFC9513, December 2023, <https://www.rfc-editor.org/info/rfc9513>.

7.2. Informative References

[RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in BGP-4", RFC 3107, DOI 10.17487/RFC3107, May 2001, <https://www.rfc-editor.org/info/rfc3107>.

[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>.

[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>.

[RFC9252] Dawra, G., Ed., Talaulikar, K., Ed., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)", RFC 9252, DOI 10.17487/RFC9252, July 2022, <https://www.rfc-editor.org/info/rfc9252>.

[RFC9855] Bashandy, A., Litkowski, S., Filsfils, C., Francois, P., Decraene, B., and D. Voyer, "Topology Independent Fast Reroute using Segment Routing", RFC 9855, DOI 10.17487/RFC9855, October 2025, <https://www.rfc-editor.org/info/rfc9855>.

Acknowledgments

The authors would like to thank Acee Lindem, Peter Psenak, Yimin Shen, Jie Dong, Zhenqiang Li, Alexander Vainshtein, Greg Mirsky, Bruno Decraene, Jeff Tantsura, Chris Bowers, Ketan Talaulikar, Bob Halley, Tal Mizrahi, Yingzhen Qu and Susan Hares for their comments to this work.

Contributors' Addresses

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

Lei Liu
Fujitsu
United States of America
Email: liulei.kddi@gmail.com

Xufeng Liu
Alef Edge
United States of America
Email: xufeng.liu.ietf@gmail.com

Authors' Addresses

Tao He
China Unicom
No.9 South Shouti Road
Beijing
100048
China
Email: het21@chinaunicom.cn

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: hchen.ietf@gmail.com

Mehmet Toy
Verizon
United States of America
Email: mehmet.toy@verizon.com

Chang Cao
China Unicom
No.9 South Shouti Road
Beijing
100048
China
Email: caoc15@chinaunicom.cn