Routing Area Working Group
Internet-Draft
Intended status: Standards Track
Expires: 8 January 2026

L. Deng
Y. Zhu
China Telecom

7 July 2025

SRv6 TE Endpoint Protection
draft-deng-rtgwg-srv6-te-endpoint-protection-00

Abstract

SRv6 TE achieves precise traffic engineering by strictly defining the traffic path (e.g., specifying particular nodes or links) through a predefined SID list. To address the failure of TI-LFA FRR protection in SRv6 TE Policy scenarios due to strict node constraints, this paper proposes an SRv6 designated Midpoint protection.

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 8 January 2026.

Table of Contents

1. Introduction

In SRv6 Best Effort (BE) forwarding, all nodes in the network rely on the same Link-State Database (LSDB) and use the Shortest Path First (SPF) algorithm to compute optimal routes, adhering to IGP shortest-path forwarding. When a link or node fails, traffic is switched to a backup path based on the IGP topology convergence. However, there may be prolonged packet loss and degraded service quality when IGP route convergence typically takes a long time in large-scale topologies. To achieve fast reroute (FRR) for SRv6 BE, a loop-free backup path must be pre-established in the forwarding plane.

The specific implementation for SRv6 BE link/node failure protection involves deploying TI-LFA (Topology-Independent Loop-Free Alternate), which enables the source node to specify an explicit path that bypasses the failed link and provide a local repair mechanism for IGP-based shortest paths. This ensures end-to-end traffic recovery in the event of direct link or neighbor failures. TI-LFA does not depend on network topology and provides a loop-free backup path. It calculates the FRR backup path based on the post-failure converged topology, ensuring alignment between the FRR backup path and the final post-convergence reroute path. This avoids secondary path switching. Specifically, TI-LFA excludes the failed node or link when computing the FRR path and then derives the backup path.

The current local repair mechanism, e.g., TI-LFA, allows the upstream neighbor of the failed node or link to fast re-route traffic around the failure. This mechanism does not work properly for SRv6 TE path after the failure happens in an endpoint node and IGP converges on the failure. This document defines midpoint protection for SRv6 TE path, which enables the upstream endpoint node of the failed node to perform the endpoint behavior for the faulty node and fast re-route traffic around the failure after IGP converges on the failure.

2. Conventions used in this document

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] .

3. Mechanism

Midpoint protection (Endpoint node) refers to the node where the SID specified in the Segment List. SRv6 Policy explicitly steers packets along a predetermined path by encapsulating a sequence of SRv6 SIDs in the Segment Routing Header (SRH). SRv6 Policy does not necessarily follow SPF shortest-path forwarding, and the IPv6 destination address is not the final packet destination. Therefore, the conventional SRv6 BE FRR which computes backup next hops based on destination addresses and SPF paths—cannot be applied in SRv6 Policy. When a link/node specified in the SRv6 Policy fails, traffic must bypass the faulty link/node via a pre-established loop-free backup path to ensure failure protection. The designated node failure protection scheme is uniformly applied in such cases.

The upstream node (Point of Local Repair, PLR) of the failed Endpoint node takes over the forwarding process. Normally, an Endpoint node processes SRv6 packets by: Decrementing the Segment Left (SL) and copying the next-layer SID to the IPv6 destination address field. However, if an Endpoint node fails, it cannot execute these SID processing actions, resulting in packet loss. Thus, when an Endpoint node fails, packets should bypass it, and the PLR must forward them directly to the next Endpoint node. When the PLR egress interface detects a failure, it executes the following protection process:

a) When the PLR’s FIB contains a route to the failed Endpoint: If the PLR is not directly connected to the failed Endpoint node, standard TI-LFA is executed. If the PLR is directly connected, it bypasses the failed Endpoint node by rewriting the IPv6 destination address to the next Endpoint node’s address in the SRH, preventing traffic from being forwarded to the faulty node.

b) When the PLR’s FIB lacks a route to the failed Endpoint: Bypass the faulty Endpoint and directly forward the data packet to the next endpoint node. If PLR detects that the next-hop interface has failed, when the next hop matches the packet’s destination address and SL > 0, the PLR proxies the Endpoint behavior: it decrements SL, updates the outer IPv6 header’s destination address with the next SID, and forwards the packet according to the new SID’s instructions, thereby bypassing the failed node and achieving SRv6 Endpoint node protection.

4. Example

As shown in Figure 1:

a) Node A forwards a packet to destination Node F, specifying intermediate Node E in the SRH.

b) When Node E fails, Node B (PLR) detects the next-hop interface failure. Since the next hop matches the current destination address 5:: and SL > 0, Node B performs proxy forwarding: Decrements SL and copies the next SID 6:: to the outer IPv6 destination field. If the segment has the Penultimate Segment Pop (PSP) attribute, node B removes the SRH and forwards based on 6::.

c) Since the primary next hop for 6:: is still Node E—but Node B is not the penultimate hop for this destination and SL = 0—Node B no longer qualifies for proxy forwarding. Instead, it follows the standard TI-LFA process, switching to a backup path. The repair segment list is encapsulated via "H.Insert" (e.g., adding SID 3::1), and a new SRH is appended before forwarding to Node F via the backup path.

d) After Node A detects Node E’ s failure and IGP convergence completes, Node A removes the routing entry for 5::. Thus, when Node A attempts to forward based on 5::, it misses the route and acts as the PLR: - Decrements SL. Updates the outer IPv6 destination address to 6::. Forwards to Node B based on 6::.

If Node B has converged, it forwards the packet to Node F via the new shortest path. If not, Node B follows TI-LFA and forwards via the backup path. To support SRv6 Endpoint failure protection, the SRv6 SID forwarding pseudocode must include additional logic to guide PLR actions during Endpoint failures.

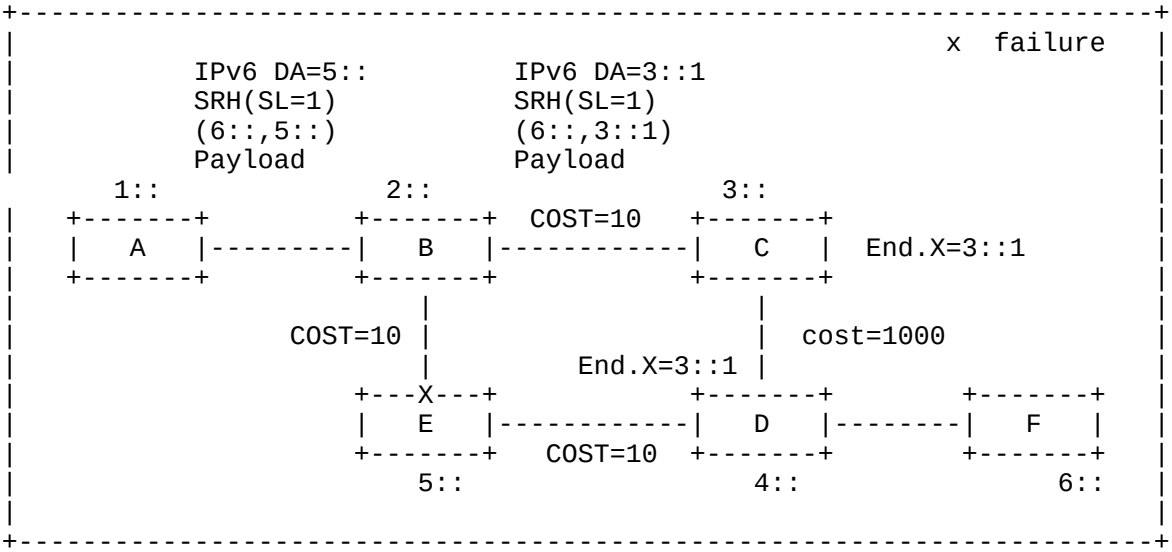

Figure 1: SRv6 Endpoint protection

5. Security Considerations

The behavior described in this document is internal functionality to a router that result in the ability to explicitly steer traffic over the post convergence path after a remote topology change in a manner that guarantees loop freeness. Because the behavior serves to minimize the disruption associated with a topology changes, it can be seen as a modest security enhancement.

6. IANA Considerations

No requirements for IANA.

7. Acknowledgement

The authors would like to thank everyone who contributed to the draft.

8. 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>.

Authors' Addresses

Lijie Deng
China Telecom
109, West Zhongshan Road, Tianhe District
Guangzhou
Guangzhou, 510000
China
Email: denglj4@chinatelecom.cn

Yongqing Zhu
China Telecom
109, West Zhongshan Road, Tianhe District
Guangzhou
Guangzhou, 510000
China
Email: zhuyq8@chinatelecom.cn