SPRING Working Group
Internet-Draft
Intended status: Standards Track
Expires: 5 January 2026

X. Geng
M. Chen
Huawei Technologies
P. Camarillo, Ed.
Cisco Systems
G. Mishra
Verizon Inc.
4 July 2025

SRv6 for Redundancy Protection
draft-ietf-spring-sr-redundancy-protection-04

Abstract

Redundancy Protection is a generalized protection mechanism to achieve high reliability of service transmission in Segment Routing networks. The mechanism uses the "Live-Live" methodology. This document introduces two new SRv6 Segment Endpoint Behavior to provide replication and elimination functions on specific network nodes by leveraging SRv6 Network Programming capabilities.

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 5 January 2026.

Copyright Notice

Table of Contents

1. Introduction 2
2. Terminology 3
2.1. Requirements Language 3
2.2. Terminology and Conventions 3
3. Redundancy Protection in Segment Routing Scenario 3
4. SRv6 Segment Endpoint Behavior to Support Redundancy Protection 5
4.1. Redundancy Segment 5
4.2. Merging Segment 6
5. Meta Data to Support Redundancy Protection 7
6. Segment Routing Policy to Support Redundancy Protection . . . 8
7. IANA Considerations 8
8. Security Considerations 9
8.1. Packet Duplication 9
8.2. Sequence Number Spoofing 9
8.3. Information Disclosure 9
8.4. State Exhaustion at Merging Node 9
9. Contributors 10
10. Acknowledgements 10
11. References 10
11.1. Normative References 10
11.2. Informative References 11
Authors' Addresses 11

1. Introduction

Redundancy Protection is a generalized protection mechanism to achieve the high reliability of service transmission in a Segment Routing (SR) network. Specifically, packets of flows are replicated at a network node into two or more copies, which are transported via different and disjoint paths in parallel. On the egress side, the multiple copies are received, redudant packets eliminated, and deliver only a single copy of the packet that is transmitted. This mechanism is commonly refered to as "Live-Live". Two new SRv6 Segment Endpoint Behavior are introduced to provide the replication and elimination functions on specific network nodes by leveraging SRv6 Network Programming capabilities. As it is unnecessary to perform switchover of recieving packets between different paths, redundancy protection can facilitate to achieve zero packet loss target when failure on either path happens.

Redundancy protection provides ultra reliable protection to many services, for example Cloud VR/Game, IPTV service and other type of video services, high value private line service etc. In this document, redundancy protection is applied to point-to-point services. The mechanism for point-to-multipoint services is out of the scope of this document.

2. Terminology

2.1. Requirements Language

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.

2.2. Terminology and Conventions

SR: Segment Routing

SRv6: Segment Routing over IPv6

SID: Segment Identifier

BSID: Binding SID

Dup node: Redundancy node. A network element that duplicates incoming packets for parallel delivery.

Mer node: Merging node. A network element that reassembles and elimantes duplicates to forward a single copy.

FID: Flow Identification

SN: Sequence Number

3. Redundancy Protection in Segment Routing Scenario

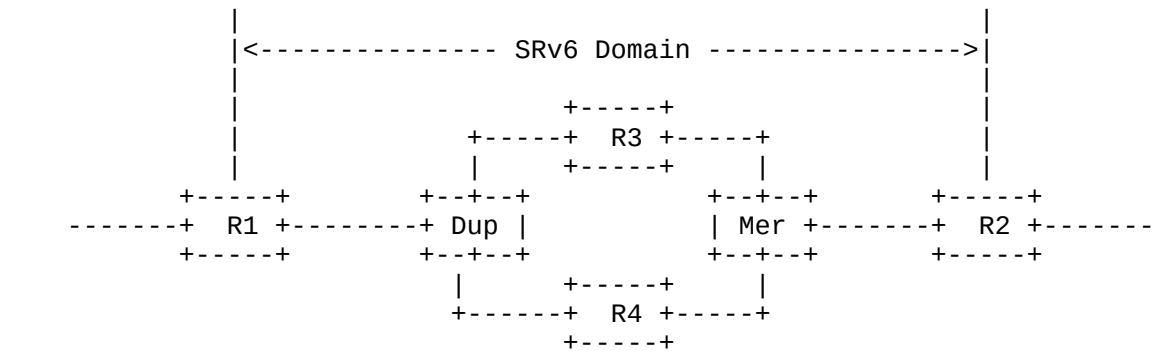

Figure 1: Example topology

Figure 1 shows an example of redundancy protection used in an SRv6 domain. R1, R2, R3, R4, Dup and Mer are SR-capable nodes. When a flow is sent into the SRv6 domain, the process is:

1) R1 receives the traffic flow and encapsulates packets with a list of segments destined to R2, which is instantiated as an ordered list of SRv6 SIDs.

2) When the packet flow arrives at Dup node, known as Redundancy node, each packet is replicated into two or more copies. Each copy of the packet is encapsulated with a new segment list, which represents different disjoint forwarding paths. The disjoint path is provisioned by a controller.

3) Meta data information including flow identification (FID) and sequence number (SN) is used to facilitate elimination of duplicate packets on Merging node (Mer). Flow identification identifies the specific flow, and sequence number distinguishes the packet sequence within a flow. This packet meta data is included on each of the replicas at the redudancy node.

4) The multiple replicas go through different paths until they reach the Mer node. The first received copy of each flow packet is transmitted from Merging node to R2, and the redundant packets are eliminated.

5) When there is any failure or packet loss in one path, the service transmission continues through the other path non-disruptively.

6) Sometimes, out-of-order packets may occur since service packets are recovered from different forwarding paths. In this case, the merging node or other network nodes behind merging node MAY include a reordering function, which is implementation specific and out of the scope of this document, to guarantee in-order delivery of packets.

To minimize the jitter caused by random packet loss, the disjoint paths are RECOMMENDED to have similar path forwarding delay.

4. SRv6 Segment Endpoint Behavior to Support Redundancy Protection

To achieve the packet replication and elimination functions, new SRv6 Endpoint Behaviors are defined.

4.1. Redundancy Segment

Redundancy Segment is the identifier of packets which need to be replicated on redundancy node. It is a variation of Binding SID (BSID) to associate with a Redundancy Policy, instantiation of which provides segment lists of different disjoint paths. Similar to the relationship between BSID and SR Policy [I-D.ietf-spring-segment-routing-policy], the use of Redundancy Segment would trigger the Redundancy Policy instantiation on redundancy node.

Redundancy Segment is associated with service instructions, indicating the following operations:

* Steers the packet into the corresponding redundancy policy

* Encapsulates flow identification and sequence number in packets if the two information is not carried in packets

* Packet replication and segment encapsulation based on the information of redundancy policy, e.g., the number of replication copies, an ordered list of segments with a topological instruction

In this document, a new behavior End.R for Redundancy Segment is defined. An instance of a redundancy SID is associated with a redundancy policy B and a source address A. In the following description, End.R behavior is specified in the encapsulation mode. The End.R behavior in the insertion mode is for further study.

When an SRv6-capable node (N) receives an IPv6 packet whose destination address matches a local IPv6 address instantiated as an SRv6 SID (S), and S is a Redundancy SID, N does:

```
S01. When an SRH is processed {
S02.   If (Segments Left>0)   {
S03.     Decrement IPv6 Hop Limit by 1
S04.     Decrement Segments Left by 1
S05.     Update IPv6 DA with Segment List[Segments Left]
S06.     Add flow identification and sequence number if indicated*
S07.     Duplicate the packets (as number of active SID lists in B)
S08.     Push the new IPv6 headers to each replica. The IPv6 header
         contains an SRH with the SID list in B
S09.     Set the outer IPv6 SA to A
S10.     Set the outer IPv6 DA to the first SID of new SRH SL
S11.     Set the outer Payload Length, Traffic Class, Flow Label,
         Hop Limit and Next-Header fields
S12.     Submit the packet to the egress IPv6 FIB lookup
         for transmission to the new destination
S13.   }
S14. }
```

* Adding flow identification and sequence number is an optional behavior for Redundancy Segment. The instruction execution is determined and explicitly indicated by SR policy or Segment itself.

4.2. Merging Segment

Merging Segment is associated with service instructions, indicates the following operations:

* Packet merging and elimination: forward the first received packets and eliminate the redundant packets

In order to eliminate the redundant packet of a flow, merging node utilizes sequence number to evaluate the redundant status of a packet. Note that implementation specific mechanism could be applied to control the amount of state monitored on sequence number, so that system memory usage can be limited at a reasonable level.

As merging node needs to maintain the state of flows, a centralized controller should have a knowledge of merging nodes capability, and never provision the redundancy policy to redundancy node when the computation result goes beyond the flow recovery capability of merging node. The capability advertisement of merging node will be specified separately elsewhere, which is not within the scope of this document.

A new SRv6 behavior for Merging Segment, End.M, is defined in this document.

When an SRv6-capable node (N) receives an IPv6 packet whose destination address matches a local IPv6 address instantiated as an SRv6 SID (S), and S is a Merging SID, N does:

```
S01. When an SRH is processed {
S02.  If (Segments Left>=0)   {
S03.    Read the SN from the packet and look it up in table
S04.      If (this sequence number does not exist in the table) {
S05.       Store this sequence number in table
S06.       Remove the outer IPv6+SRH header
S07.       Decrement IPv6 Hop Limit by 1 in inner SRH
S08.       Decrement Segments Left by 1 in inner SRH
S09.       Update IPv6 DA with Segment List[Segments Left] in inner SRH
S10.       Submit the packet to the egress IPv6 FIB lookup and transmit
S11.      }
S12.      ELSE {
S13.           Drop the packet
S14.      }
S15.    }
S16. }
```

5. Meta Data to Support Redundancy Protection

To support the redundancy protection function, flow identification and sequence number are added in the packet and further used at merging node when the merging function is executed. Flow identification identifies one specific flow of redundancy protection, and is usually allocated from centralized controller to SR ingress node or redundancy node in SR network. Note that flow identification can also be allocated and advertised by merging node. BGP, PCEP or Netconf protocols can facilitate the advertisement and distribution of flow identification among controller, redundancy node and merging node. Sequence number distinguishes the packets within a flow by specifying the order of packets. Not like the uniqueness of flow identification to one specific flow, sequence number keeps changing to each packet within a flow. It is RECOMMENDED to add the sequence number in forwarding plane as performance and scalability is required.

Figure 2 suggests an encapsulation of flow identification and sequence number in Segment Routing Header (SRH)[RFC8754] when redundancy protection is used in SRv6 network.

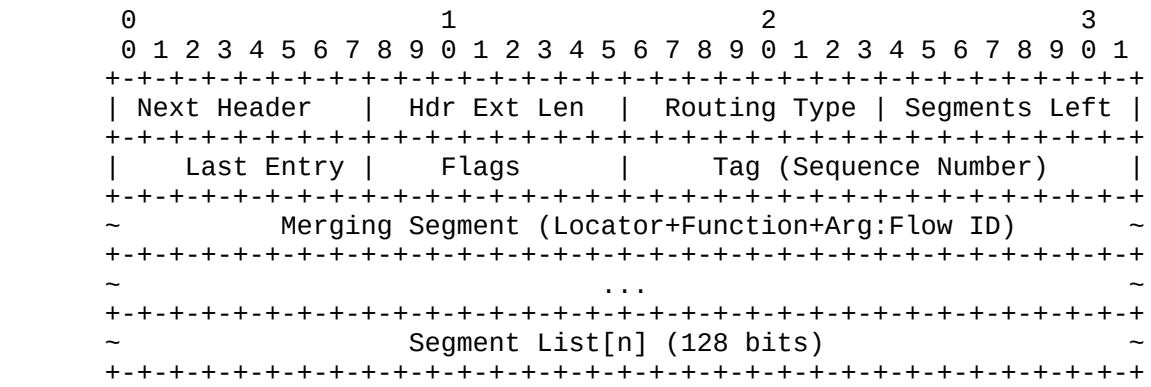

Figure 2 Encapsulation of Flow Identification and Sequence Number

Since the flow identification is only used at merging node to identify the specific flow of redundancy protection, it is encapsulated in the Arguments of Merging Segment in SRH. The length of flow identification is not limited, however in practice it is suggested to be 16 bits.

All the duplicates of the same packet need to be tagged for deduplication at the merging node. For this purpose, we will use a sequence number. It is RECOMMENDED to encode the seq number in the Tag field of the SRH, with a length of 16bits.

6. Segment Routing Policy to Support Redundancy Protection

Redundancy Policy is a variation of SR Policy to conduct the replicas to multiple disjoint paths for redundancy protection. It extends SR policy [I-D.ietf-spring-segment-routing-policy] to include more than one active and parallel ordered lists of segments between redundancy node and merging node, and all the ordered lists of segments are used at the same time to steer each copy of flow into different disjoint paths.

7. IANA Considerations

This document requires registration of End.R behavior and End.M behavior in "SRv6 Endpoint Behaviors" sub-registry of "Segment Routing Parameters" registry.

IANA maintains The "SRv6 Endpoint Behaviors" sub-registry of the "Segment Routing Parameters" registry. IANA is reuqested to make two new assignments from the First Come First Served portion of the registry as follows:

Value	Hex	Endpoint Behavior	Reference	Change Controller
TBD1	xTBD1	End.R	[This.I-D]	IETF
TBD1	xTBD1	End.M	[This.I-D]	IETF

8. Security Considerations

The introduction of Redundancy Segments and Merging Segments in Segment Routing networks introduces new vectors for security threats that must be carefully mitigated.

8.1. Packet Duplication

Redundancy protection intentionally replicates packets across multiple paths. Without proper admission control or policy enforcement, an attacker could exploit this mechanism to amplify traffic, overwhelming downstream links or merging nodes.

The use of redundancy protection SHOULD be restricted to trusted applications and provisioned via authenticated and authorized controllers (e.g., using BGP with RPKI or PCEP with TLS). Rate-limiting and flow admission control at the ingress SHOULD be employed to prevent abuse.

8.2. Sequence Number Spoofing

The merging node relies on sequence numbers to de-duplicate packets. An attacker that can inject or manipulate these sequence numbers could cause legitimate packets to be dropped or reordered.

Redundancy Segments MUST be deployed only within trusted SR domains.

8.3. Information Disclosure

Redundancy protection may involve topology-specific path selections that reveal operational characteristics of the network (e.g., availability of disjoint paths).

Such information SHOULD NOT be exposed outside the trusted SR domain. Control-plane interactions involving Redundancy Segments SHOULD be encrypted and authenticated (e.g., BGP with TCP-AO, PCEP over TLS).

8.4. State Exhaustion at Merging Node

Merging nodes need to maintain state (e.g., sequence windows, buffering) for each redundancy-protected flow. An attacker might attempt to create many such flows to exhaust memory or processing capacity.

Merging nodes SHOULD limit the number of concurrent redundancy flows per source. Idle timeout mechanisms MUST be implemented to garbage-collect stale state.

9. Contributors

Fan Yang
Huawei
China
Email: shirley.yangfan@huawei.com

10. Acknowledgements

The authors would like to thank Bruno Decraene, Ron Bonica, James Guichard, Jeffrey Zhang, Balazs Varga, Adrian Farrel for their valuable comments and discussions.

11. References

11.1. Normative References

[I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", Work in Progress, Internet-Draft, draft-ietf-spring-segment-routing-policy-22, 22 March 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-segment-routing-policy-22>.

[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>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <https://www.rfc-editor.org/info/rfc8754>.

[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>.

11.2. Informative References

[RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <https://www.rfc-editor.org/info/rfc8655>.

Authors' Addresses

Xuesong Geng
Huawei Technologies
China
Email: gengxuesong@huawei.com

Mach(Guoyi) Chen
Huawei Technologies
China
Email: mach.chen@huawei.com

Pablo Camarillo Garvia (editor)
Cisco Systems
Spain
Email: pcamaril@cisco.com

Gyan Mishra
Verizon Inc.
Email: gyan.s.mishra@verizon.com