Workgroup: SPRING
Internet-Draft:
draft-geng-spring-redundancy-policy-05
Published: 13 March 2023
Intended Status: Standards Track
Expires: 14 September 2023
Authors: F. Yang X. Geng T. Zhou G. Mishra
Huawei Huawei Huawei Verizon Inc.

Redundancy Policy for Redundancy Protection

Abstract

This document introduces a variant of SR Policy called Redundancy Policy, in order to instruct the replication of service packets and assign more than one redundancy forwarding paths used for redundancy protection.

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 .

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

Table of Contents

1. Introduction

Redundancy protection [I-D.ietf-spring-sr-redundancy-protection] is a generalized protection mechanism by replicating and transmitting copies of flow packets on redundancy node over multiple different and disjoint paths, and further eliminating the redundant packets at merging node. This document introduces Redundancy Policy to support redundancy protection, which is a variant of SR Policy [I-D.ietf-spring-segment-routing-policy]. Redundancy Policy instructs the replication of service packets and assigns more than one equivalent forwarding paths used for redundancy protection. Redundancy Policy applies equally to both MPLS data plane (SR-MPLS) [RFC8660] and Segment Routing with IPv6 data plane (SRv6) [RFC8986].

2. Terminology and Conventions

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.

The other terminologies used in this document are:

Redundancy Node: the start point of redundancy protection, where the network node replicates the flow packets.

Merging Node: the end point of redundancy protection, where the network node eliminates and ordering(optionally) the flow packets.

Redundancy Policy: an extended SR Policy which instructs more than one redundancy forwarding paths to support packet redundant transmission.

3. Redundancy Policy

Redundancy Policy is used to enable packet replication and instantiation more than one active ordered lists of segments between redundancy node and merging node to steer the same flow through different paths in an SR domain.

3.1. Identification of Redundancy Policy

Redundancy Policy is a variant of SR Policy and also identified through the tuple <headend, color, endpoint>. Specifically, a redundancy policy is identified by <redundancy node, color, merging node>. Redundancy node is specified as IPv4/IPv6 address of headend of Redundancy Policy, which is the node to perform packet replication. Merging node is specified as IPv4/IPv6 address of endpoint of Redundancy Policy, which is the node to perform packet elimination. The value of color specifies the intent of the redundancy policy is "redundancy protection for high reliability", which indicates service packets are replicated into multiple copies and carried on different forwarding paths .

3.2. Structure of Redundancy Policy

Redundancy policy shares the basic structure and elements with SR Policy and its information model is shown in the following:

```
Redundancy policy POL1 <R Node= R1, Color = 1, M Node = M1>
        Candidate-path CP1 <protocol-origin = 20, originator
                           = 100:1.1.1.1, discriminator = 1>
             Flag Redundancy
             Preference 200
             SID-List1 <SID11...SID1i>
             SID-List2 <SID21...SID2j>
        Candidate-path CP2 <protocol-origin = 20, originator
                           = 100:2.2.2.2, discriminator = 2>
             Preference 100
             Weight W3, SID-List3 <SID31...SID3i>
```

The Redundancy Policy POL1 is identified by the tuple <redundancy node, color, merging node>, in which R1 is the redundancy node, M1

is the merging node, and Color 1 represents the intent of redundancy protection. Two candidate-paths CP1 and CP2 instruct the ordered segment lists from redundancy node to merging node. In candidate path CP1, a new attribute Flag is added to indicate the type of candidate path. When the candidate path is indicated with the flag of redundancy, the attribute Weight is not applicable to the SID-Lists and all SID Lists of the candidate path are used for redundancy forwarding. Regarding the other attributes of candidate path such as originator, preference, priority, segment-list etc, the definitions apply the same as [I-D.ietf-spring-segment-routing-policy].

3.3. Flag of a Candidate Path

Flag is an optional attribute of a candidate path, which is used to indicate the type of a candidate path is for redundancy forwarding. When the candidate path with flag of redundancy is selected as the active candidate path, this SR Policy is identified as the Redundancy Policy. Flag of a candidate path is an 8-bit bitmap. The table below specifies the current definition of Flag:

Bitmap	Flag	Description
0	R	Redundancy paths
1-7	U	Reserved

Figure 2: Flag

3.4. Behavior of Redundancy Policy

When the SR policy is identified as a redundancy policy, network node uses rules to compute and select the valid active ordered segment-lists for redundancy forwarding. The specific rules are:

- The candidate paths are selected to determine the best path of an SR policy. Preference, Protocol-Origin, and other tie-breaking rules defined in section 2.9 of [I-D.ietf-spring-segment-routing-policy] are evaluated until only one valid best path is selected.

- In a redundancy policy, the candidate path with a flag of redundancy is always selected as the best path in the first place.

- When the selected active candidate path is with a flag of redundancy, all the segment-lists of the candidate path are used

as the active segment-lists for redundancy forwarding, where each active segment-list carries an entire copy of service packets.

*Weight is not applicable for the segment-lists in a candidate path with a flag of redundancy. Redundancy policy has no purpose of weighted load-balancing.

*The candidate path without a flag of redundancy in the same SR policy with the candidate paths with a flag, is considered as the backup path, which allowing provisioning of multiple path options.

Take the information model in section 3.2 as an example, preference value 200 of CP1 is higher than preference value 100 of CP2, thus CP1 is selected as the active candidate path. Because CP1 is with the flag of redundancy, both Segment-List1 and Segment-List2 are selected as the active Segment-Lists for redundancy forwarding. After service packets are replicated, each segment-list forwards each replicas of service packets. When CP1 becomes invalid and fallbacks to CP2, CP2 provides the backup path to the redundancy forwarding.

3.5. BSID and Redundancy Policy

Redundancy policy can be optionally associated with a Binding Segment. Redundancy SID defined in [I-D.ietf-spring-sr-redundancy-protection] can be the Binding SID of redundancy policy. In other words, Redundancy SID triggers the instantiation of redundancy policy in the forwarding plane on redundancy node.

3.6. Steering into a Redundancy Policy

A packet is steered into a Redundancy Policy at a redundancy node in following ways:

*Incoming packets have an active SID matching the Redundancy SID at the redundancy node.

*Per-destination Steering: incoming packets match a BGP/Service route which recurses on a Redundancy Policy.

*Per-flow Steering: incoming packets match or recurse on a forwarding array of where some of the entries are Redundancy Policy.

*Policy-based Steering: incoming packets match a routing policy which redirects them on a Redundancy Policy.

3.7. Protocol Extensions

Similar to SR Policy, Redundancy Policy requires the control plane protocol extensions to distribute candidate paths and other information. New sub-TLVs are expected to be defined to encode new information of Redundancy Policy Candidate Paths in BGP [I-D.ietf-idr-segment-routing-te-policy] and PCEP [I-D.ietf-pce-segment-routing-policy-cp].

4. IANA Considerations

TBD

5. Security Considerations

TBD

6. References

6.1. 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>.

[RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019, <https://www.rfc-editor.org/info/rfc8660>.

[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>.

6.2. Informative References

[I-D.ietf-idr-segment-routing-te-policy]
Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-segment-routing-te-policy-20, 27 July 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-segment-routing-te-policy-20>.

[I-D.ietf-pce-segment-routing-policy-cp]
Koldychev, M., Sivabalan, S., Barth, C., Peng, S., and H. Bidgoli, "PCEP extension to support Segment Routing Policy Candidate Paths", Work in Progress, Internet-Draft, draft-ietf-pce-segment-routing-policy-cp-09, 7 March 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-pce-segment-routing-policy-cp-09>.

[I-D.ietf-spring-sr-redundancy-protection]
Geng, X., Chen, M., Yang, F., Camarillo, P., and G. S. Mishra, "SRv6 for Redundancy Protection", Work in Progress, Internet-Draft, draft-ietf-spring-sr-redundancy-protection-02, 23 September 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-sr-redundancy-protection-02>.

Authors' Addresses

Fan Yang
Huawei
156 Beiqing Rd.
Beijing
100095
China

Email: shirley.yangfan@huawei.com

Xuesong Geng
Huawei
156 Beiqing Rd.
Beijing
100095
China

Email: gengxuesong@huawei.com

Tianran Zhou
Huawei
156 Beiqing Rd.
Beijing

100095
China

Email: zhoutianran@huawei.com

Gyan Mishra
Verizon Inc.

Email: gyan.s.mishra@verizon.com