

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 8 September 2022

X. Geng
M. Chen
F. Yang, Ed.
Huawei Technologies
7 March 2022

Redundancy Policy for Redundancy Protection
draft-geng-spring-redundancy-policy-02

Abstract

Redundancy Protection is a generalized protection mechanism to achieve the high reliability of service transmission in Segment Routing 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. To support redundancy protection in Segment Routing domain, this document introduces Redundancy Policy, as a variant of SR Policy, to intrust the replication of service packets and the multiple ordered lists of segments used for packet carrying.

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 8 September 2022.

Internet-Draft

Redundancy Policy

March 2022

Copyright Notice

Copyright (c) 2022 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
2.	Terminology and Conventions	2
3.	Redundancy Policy	3
3.1.	Identification of Redundancy Policy	3
3.2.	Redundancy Policy Structure	4
3.3.	Behavior of Redundancy Policy	4
3.4.	Association with Redundancy Segment	5
4.	IANA Considerations	5
5.	Security Considerations	5
6.	Acknowledgements	5
7.	References	5
7.1.	Normative References	5
7.2.	Informative References	6
	Authors' Addresses	7

[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 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", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

Geng, et al.

Expires 8 September 2022

[Page 2]

Internet-Draft

Redundancy Policy

March 2022

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 active segment lists for the multiple paths to support redundant transmission of flow packets.

[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

A Redundancy Policy is identified through the tuple <redundancy node, redundancy ID, 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. Redundancy ID could be a specified value of "color" define in section 2.1 of [[I-D.ietf-spring-segment-routing-policy](#)], which indicates the SR policy as a redundancy policy. Redundancy ID could also be used to distinguish redundancy policy sharing the same redundancy node and merging node.

The following elements are extended in Redundancy Policy:

- * **Redundancy ID:** is used to distinguish a redundancy policy or different redundancy policies

- * Redundancy SID: is the Binding SID for a redundancy policy and defined in [[I-D.ietf-spring-sr-redundancy-protection](#)]. Redundancy SID triggers the instantiation of a Redundancy Policy in the forwarding plane of redundancy node.
- * Candidate path: more than one candidate paths are included in redundancy policy. In each candidate path, the last segment SHOULD be the Merging SID defined in [[I-D.ietf-spring-sr-redundancy-protection](#)]. The preference of candidate paths in redundancy policy SHOULD be the same to instantiate more than one active ordered segment lists.

[3.2.](#) Redundancy Policy Structure

Redundancy policy inherits the basic structure and elements of SR Policy, the information model of Redundancy Policy is the following:

```

Redundancy policy POL1 <R Node= R1, Redundancy ID = 1, M Node = M1>
  Candidate-path CP1 <originator = 100:1.1.1.1>
    Preference 200
    Priority 10
    Weight W1, SID-List1 <SID11...SID1i>
    Weight W2, SID-List2 <SID21...SID2j>
  Candidate-path CP2 <originator = 100:1.1.1.1>
    Preference 200
    Priority 10
    Weight W3, SID-List3 <SID31...SID3i>

```

The Redundancy Policy POL1 is identified by the tuple <redundancy node, redundancy ID, merging node>, R1 is the redundancy node, and M1 is the merging node. Redundancy ID is used to identify as a redundancy policy in the context of redundancy node. Two candidate-paths CP1 and CP2 instruct the ordered segment lists, with the same definition of SR policy. Preference is mandatory for each candidate path and used to determine the parallel forwarding paths. Candidate paths with the same preference value are chosen as the disjoint forwarding path of redundancy protection. Since there is no tie-breaking rules of the only one valid best path selection, attributes of protocol-origin and discriminator are not applicable in redundancy policy.

[3.3.](#) Behavior of Redundancy Policy

In Redundancy Policy, the preference of candidate path is used to select the valid active candidate paths. The preference of candidate paths in redundancy policy SHOULD be the same. Different from SR Policy, there is no tie-breaker comparison between candidate path with equal preference values. Redundancy Policy has no limit on the number of active CPs since more than one forwarding paths are required in order to perform redundancy protection. Regarding the instantiation of ordered lists of segments for candidate path, it keeps the same forwarding instantiation and behavior defined for SR policy. For example, traffic steered on POL1 is still flow-based hashed on Segment-List1 <SID11...SID1i> with a ratio $W1/(W1+W2)$, so does Segment-List2.

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 directs them on a Redundancy Policy.

[3.4.](#) Association with Redundancy Segment

In Redundancy Policy, each candidate path must be associated with a BSID. The endpoint behavior of BSID specifies the instantiation of Redundancy Policy in forwarding plane and adopts the endpoint behavior definition of Redundancy SID

[[I-D.ietf-spring-sr-redundancy-protection](#)] . The associated BSID of Redundancy Policy and its endpoint behavior are signalled redundancy node via BGP SR Policy [[I-D.ietf-idr-segment-routing-te-policy](#)] or

PCEP [[RFC8664](#)] [[I-D.ietf-pce-segment-routing-policy-cp](#)] or Netconf YANG.

[4.](#) IANA Considerations

This document requires no new registration IANA.

[5.](#) Security Considerations

TBD

[6.](#) Acknowledgements

Thanks for the valuable comments from James Guichard and Andrew Mail.

[7.](#) References

[7.1.](#) Normative References

Geng, et al.

Expires 8 September 2022

[Page 5]

Internet-Draft

Redundancy Policy

March 2022

[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-18](#), 17 February 2022, <<https://www.ietf.org/archive/id/draft-ietf-spring-segment-routing-policy-18.txt>>.

[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>>.
- [RFC8964] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., Bryant, S., and J. Korhonen, "Deterministic Networking (DetNet) Data Plane: MPLS", [RFC 8964](#), DOI 10.17487/RFC8964, January 2021, <<https://www.rfc-editor.org/info/rfc8964>>.
- [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>>.

7.2. 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-14](#), 10 November 2021, <<https://www.ietf.org/archive/id/draft-ietf-idr-segment-routing-te-policy-14.txt>>.
- [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-06, 22 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-pce-segment-routing-policy-cp-06.txt>>.

- [I-D.ietf-spring-sr-policy-yang]
Raza, K., Sawaya, R., Shunwan, Z., Voyer, D., Durrani, M., Matsushima, S., and V. P. Beeram, "YANG Data Model for Segment Routing Policy", Work in Progress, Internet-Draft,

[draft-ietf-spring-sr-policy-yang-01](https://www.ietf.org/archive/id/draft-ietf-spring-sr-policy-yang-01), 7 April 2021,
<<https://www.ietf.org/archive/id/draft-ietf-spring-sr-policy-yang-01.txt>>.

[I-D.ietf-spring-sr-redundancy-protection]

Geng, X., Chen, M., Yang, F., Garvia, P. C., and G. Mishra, "SRv6 for Redundancy Protection", Work in Progress, Internet-Draft, [draft-ietf-spring-sr-redundancy-protection-01](https://www.ietf.org/archive/id/draft-ietf-spring-sr-redundancy-protection-01), 15 February 2022,
<<https://www.ietf.org/archive/id/draft-ietf-spring-sr-redundancy-protection-01.txt>>.

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

[RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", [RFC 8664](https://www.rfc-editor.org/info/rfc8664), DOI 10.17487/RFC8664, December 2019,
<<https://www.rfc-editor.org/info/rfc8664>>.

Authors' Addresses

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

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

Fan Yang
Huawei Technologies
China

