SPRING Working Group
Internet Draft
Intended status: Standards Track
Expires: December 30, 2025

S. Yue
China Mobile
R. Chen
ZTE Corporation
J. Dong
Huawei Technologies
C. Lin
New H3C Technologies
R. Pang
China Unicom
K. Zhang
Huawei Technologies
W. Gao
CAICT
W. Jiang
China Mobile

June 30, 2025

Segment Routing Policy Extension for Network Resource Partition

draft-jiang-spring-sr-policy-nrp-03

Abstract

Segment Routing (SR) Policy is a set of candidate paths, each consisting of one or more segment lists and the associated information. A Network Resource Partition (NRP), is a subset of the resources and associated policies in the underlay network. In SR networks with multiple NRPs, an SR Policy can be associated with a particular NRP. In that case, SR Policy can be used for steering and forwarding traffic which is mapped to the NRP, so that the packets can be processed with the subset of network resources and policy of the NRP for guaranteed performance. Thus the association between SR Policy and NRP needs to be specified.

This document describes how the SR Policy extension for associated NRP and the operational mechanisms function together.

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 December 30, 2025.

Table of Contents

1. Introduction

A Segment Routing Policy (SR Policy) [RFC9256] is a set of candidate paths, each consisting of one or more segment lists and the associated information. The headend node is said to steer a flow into an SR Policy. The packets steered into an SR Policy have an ordered list of segments associated with that SR Policy written into them. [RFC8660] describes the representation and processing of this ordered list of segments as an MPLS label stack for SR-MPLS, while [RFC8754] and [RFC8986] describe the same for Segment Routing over IPv6 (SRv6) with the use of the Segment Routing Header (SRH).

[RFC9543] provides the definition of IETF network slice for use within the IETF and discusses the general framework for requesting and operating IETF Network Slices, their characteristics, and the necessary system components and interfaces.It also introduces the concept Network Resource Partition (NRP), which is a subset of the resources and associated policies in the underlay network.

In SR networks, an NRP can be realized using NRP-specific resource-aware segments as defined in [I-D.ietf-spring-resource-aware-segments]. With this approach, for each NRP, a separate set of resource-aware SIDs need to be assigned, thus the amount of SR SIDs would be proportional to the number of NRPs.

As described in [I-D.ietf-teas-nrp-scalability], one scalable data plane approach to support network slicing is to carry a dedicated NRP Selector ID in the data packet to identify the NRP the packet belongs to, so that the packet can be processed and forwarded using the subset of network resources allocated to the NRP.

In SR networks with multiple NRPs, an SR Policy can be associated with a particular NRP. In that case, SR Policy can be used for steering and forwarding traffic which is mapped to the NRP, so that the packets can be processed with the subset of network resources and policy of the NRP for guaranteed performance. Thus the association between SR Policy and NRP needs to be specified.

This document describes how the SR Policy extension for associated NRP and the operational mechanisms function together.

1.1. 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. Use Case

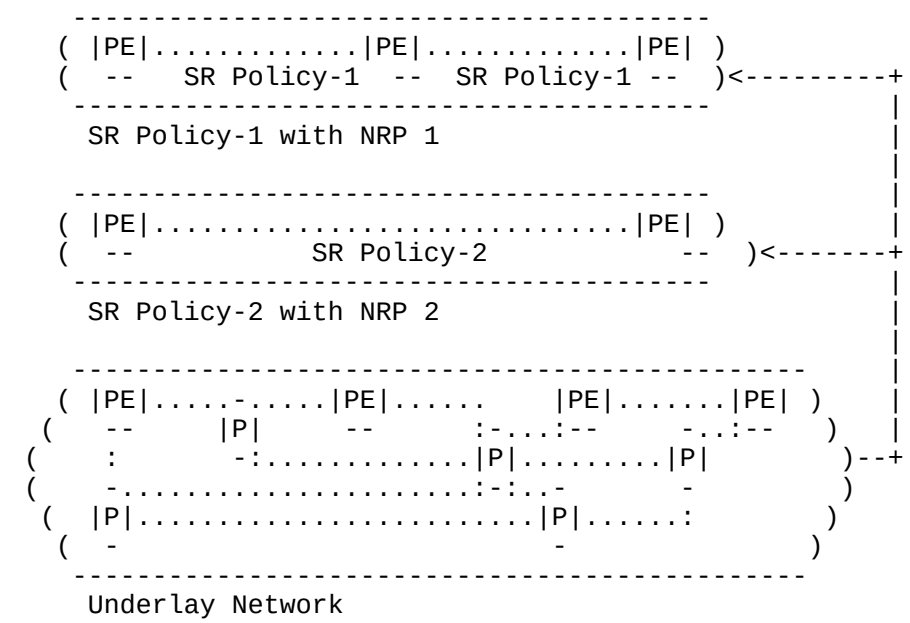

Figure 1: Use case of SR Policy Extension for NRP

In each NRP for network slices, the connectivity among PEs is achieved by SR Policies. The segment lists of these SR Policies composed with segments associated with the dedicated data plane NRP Selector ID. Traffics are steered into the SR Policies, so that the resources allocated to the corresponding NRPs will be used for forwarding.

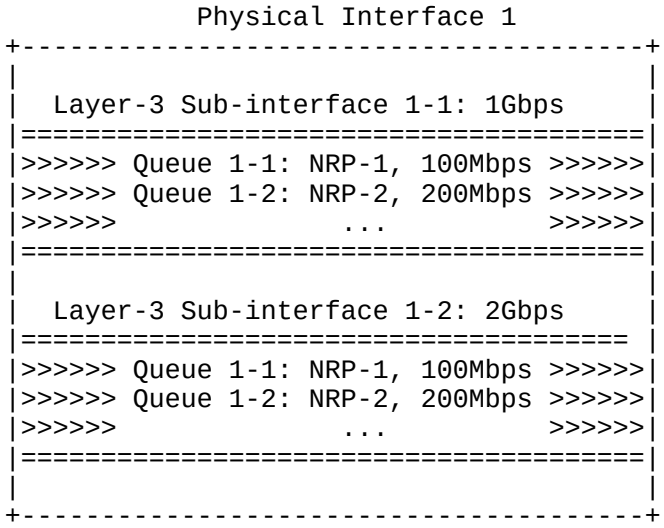

Figure 2: Network Resource Partition on An Interface

As shown in the example in Figure 2, the bandwidth resource of a physical interface is partitioned in two NRPs.

The NRPs are sliced by HQoS queues with dedicated bandwidth under the layer-3 sub-interface. NRP needs to be identified by using an extra dimension. On both MPLS-SR and SRv6 data plane, there are several options for realizing NRP Selector ID, such as [I-D.ietf-6man-enhanced-vpn-vtn-id], [I-D.cheng-spring-srv6-encoding-network-sliceid], and [I-D.li-mpls-enhanced-vpn-vtn-id]. As mentioned above, the traffics of network slice are forwarded according to the segment list of SR Policy. Firstly, the outgoing interface associated segment will be the layer-3 sub-interface. Then, the HQoS queue will be selected according to the NRP Selector ID carried in the packets, and the bandwidth resource of NRP will be used.

3. SR Policy Extension for NRP

As defined in [RFC9256], an SR Policy is associated with one or more candidate paths. A candidate path is the unit for signaling of an SR Policy to a headend via protocol extensions like the Path Computation Element Communication Protocol (PCEP) [RFC8664] [I-D.ietf-pce-segment-routing-policy-cp] or BGP SR Policy [I-D.ietf-idr-sr-policy-safi]. A candidate path consists of one or multiple segment lists. The segment lists are used for load balancing purpose. When an SR Policy is associated with an NRP, the SR Policy is instantiated using candidate paths which are built within a particular NRP. Hence the association between SR Policy and NRP is

specified at the candidate path level. All the segment lists of the candidate path are associated with the same NRP and share the set of resources of the NRP.

The candidate paths of an SR Policy determine the path that packets will traverse, while NRP reserves resources along the candidate path designated by the SR Policy. Through the integration of SR Policy and NRP, it ensures both the forwarding path and resource reservation along the candidate path.

3.1. NRP Selector ID of a Candidate Path

The NRP Selector ID of a candidate path is utilized to identify the resources corresponding to the forwarding paths of all segment lists within an SR Policy. It is a 32-bit value serving as an identifier for the Network Resource Partition. The NRP Selector ID associated with a candidate path of an SR Policy from a specific Protocol-Origin as specified below:

- When provisioning is via configuration, it is specific to the implementation’s configuration model.
- When signaling is via PCEP, the method to uniquely signal an individual candidate path along with its NRP Selector ID is described in [I-D.draft-dong-pce-pcep-nrp].
- When signaling is via BGP SR Policy, the method to uniquely signal an individual candidate path along with its NRP Selector ID is described in [I-D.ietf-idr-sr-policy-nrp].It can be collected via BGP-LS [I-D.draft-ietf-idr-bgp-ls-sr-policy-nrp].

Under the same Candidate Path, all segment lists must share the same NRP Selector ID. When a candidate path of an SR Policy is instantiated within an NRP, a network-wide data plane NRP Selector ID is used to identify the resources of the NRP. While different candidate paths can share the same NRP Selector IDs, the proposed mechanism allows for different candidate paths within a single SR Policy to be associated with different NRPs. However, in typical network scenarios, it is generally expected that the association between an SR Policy and an NRP remains consistent. In such cases, all candidate paths of a single SR Policy SHOULD be associated with the same NRP.

By associating NRP Selector IDs with Candidate Paths, the assurance of both the SR Policy’s path and its resources is achieved. The process involves the following steps:

- Planning the network topology resources and assigning NRP Selector IDs.

- At the headend node, performing path arrangement. During the path planning process of the SR Policy, resources are considered for different candidate paths, and NRP Selector IDs are configured under each Candidate path to establish the association between the path and resources.

3.2. Candidate Path Validity Verification

A candidate path is considered usable when it is valid, with the validation rules outlined in Section 5 of [RFC9256]. When a Candidate Path contains an NRP Selector ID, a segment list of a candidate path may be declared invalid if the resources corresponding to the NRP Selector ID on the segment list path do not exist. The successful reservation of NRP resources along the entire path can be verified through OAM (Operations, Administration, and Maintenance) detection mechanisms. Additionally, if the head-end is unable to perform path resolution for the first SID into one or more outgoing interfaces and next-hops, along with the corresponding NRP Selector ID resources, the status of that segment list is set to invalid.

When running fast detection protocols, such as Bidirectional Forwarding Detection (BFD), the headend may compute and validate backup candidate paths and provision them into the forwarding plane as a backup for the active path. In such cases, it is necessary to include NRP encapsulation to detect the NRP resources along the path, ensuring the availability of both the path and resources.

3.3. Summary

For an SR Policy associated with an NRP, each of its candidate paths must be associated with an NRP. The NRP Selector ID linked to each candidate path can be the same or different. All segment lists of the candidate path are associated with the same NRP and share the set of resources allocated to that NRP.

In summary, the information model is the following:

SR Policy POL1

Candidate Path CP1

Preference 200

NRP Selector ID 100

```
            Segment List 1 <SID11...SID1i>, Weight 1

            Segment List 2 <SID21...SID2j>, Weight 1

            Segment List 3 <SID31...SID3k>, Weight 1

         Candidate Path CP2

         Preference 100

         NRP Selector ID 200

            Segment List 4 <SID41...SID4i>, Weight 1

            Segment List 5 <SID51...SID5j>, Weight 1

            Segment List 6 <SID61...SID6k>, Weight 1
```

SR Policy POL1 has two Candidate Paths, CP1 and CP2. CP1 is the active candidate path (valid and with the highest Preference). NRP Selector ID 100 is configured under CP1, while NRP Selector ID 200 is configured under CP2. The three segment lists of CP1 with NRP Selector ID 100 are installed as the forwarding instantiation of SR Policy POL1. NRP Selector ID 100 needs to be configured and resources reserved on the paths traversed by segment list 1, segment list 2, and segment list 3. When traffic is steered on POL1 and flow-based hashed on segment list <SID11...SID1i>, NRP-100 is added into the data packet, and forwarding is based on the resources pointed to by NRP-100.

4. Steering into an SR Policy with NRP

The method of traffic steering aligns with the description in Section 8 of [RFC9256]. If the SR Policy candidate path selected as the best candidate path is associated with an NRP, the headend node of the SR Policy SHOULD encapsulate both the segment list and the data plane identifier of the associated NRP Selector ID to the header of packets steered to the SR Policy. The segment list is used to instruct the path the packets need to traverse, and the NRP Selector ID is used by each node along the path to identify the set of local network resources allocated to the NRP for the processing of the packet.When an SR policy’s active path contains an NRP Selector ID, specific handling is necessary, as follows:

- When steering traffic to the SR policy through Per-Destination Steering or Policy-Based Routing, after adding the corresponding segment list encapsulation for the SR policy, NRP encapsulation is also required. The specific NRP encapsulation details are outside the scope of this document.

- Similarly, When steering traffic to the SR policy via the BindingSID, after adding the segment list encapsulation for the SR policy, NRP encapsulation is required. The specific NRP encapsulation details are outside the scope of this document.

5. Manageability Considerations

This document specifies the detailed construction of the SR Policy with NRP Selector ID and its operational mechanisms. Therefore, the manageability considerations of [RFC9256] apply.

6. Security Considerations

The security considerations described in [RFC9256] also apply to this document.

This document does not introduce any new security consideration.

7. IANA Considerations

This document has no IANA actions.

8. References

8.1. 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>.

[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>.

[RFC9256] Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", BCP 14, RFC 9256, DOI 10.17487/RFC9256, July 2022, < https://www.rfc-editor.org/info/rfc9256 >.

[RFC9543] Farrel, A., Ed., Drake, J., Ed., Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "A Framework for Network Slices in Networks Built from IETF Technologies", RFC 9543, DOI 10.17487/RFC9543, March 2024, <https://www.rfc-editor.org/info/rfc9543>.

8.2. Informative References

[I-D.ietf-teas-nrp-scalability] Dong, J., Li, Z., Gong, L., Yang, G., Mishra, G. S., and F. Qin, "Scalability Considerations for Network Resource Partition", Work in Progress, Internet-Draft, draft-ietf-teas-nrp-scalability-03, 21 October 2023,<https://datatracker.ietf.org/doc/html/draft-ietf-teas-nrp-scalability-03>.

[I-D.ietf-6man-enhanced-vpn-vtn-id] Dong, J., Li, Z., Xie, C., Ma, C., and G. Mishra, "Carrying Virtual Transport Network (VTN) Identifier in IPv6 Extension Header", Work in Progress, Internet-Draft, draft-ietf-6man-enhanced-vpn-vtn-id-00, 5 March 2022, <http://www.ietf.org/internet-drafts/draft-ietf-6man-enhanced-vpn-vtn-id-00.txt>.

[I-D.cheng-spring-srv6-encoding-network-sliceid] Cheng, W., Lin, C., Gong, L., Zadok, S., and X. Wang, "Encoding Network Slice Identification for SRv6", Work in Progress, Internet-Draft, draft-cheng-spring-srv6-encoding-network-sliceid-04, 8 July 2022, <http://www.ietf.org/internet-drafts/draft-cheng-spring-srv6-encoding-network-sliceid-04.txt>.

[I-D.decraene-mpls-slid-encoded-entropy-label-id] Decraene B., Filsfils, C., Henderickx W., Saad T., Beeram V., "Using Entropy Label for Network Slice Identification in MPLS networks", Work in Progress, Internet-Draft, draft-decraene-mpls-slid-encoded-entropy-label-id-04, 14 June 2022, <http://www.ietf.org/internet-drafts/draft-decraene-mpls-slid-encoded-entropy-label-id-04.txt>.

[I-D.li-mpls-enhanced-vpn-vtn-id] Li, Z. and J. Dong, "Carrying Virtual Transport Network Identifier in MPLS Packet", Work in Progress, Internet-Draft, draft-li-mpls-enhanced-vpn-vtn-id-02, 7 March 2022, <http://www.ietf.org/internet-drafts/draft-li-mpls-enhanced-vpn-vtn-id-02.txt>.

[I-D.ietf-pce-segment-routing-policy-cp] Koldychev, M., Sivabalan, S., Barth, C., Peng, S., and H. Bidgoli, "Path Computation Element Communication Protocol PCEP) Extensions for Segment Routing (SR) Policy Candidate Paths", Work in Progress, Internet-Draft, draft-ietf-pce-segment-routing-policy-cp-18, 14 October 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-pce-segment-routing-policy-cp-18>.

[I-D.ietf-idr-sr-policy-safi] Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-safi-09, 3 October 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-safi-09>.

[I-D. draft-dong-pce-pcep-nrp] Dong, J.,S. Fang,Q. Xiong,S. Peng,L. Han, "Path Computation Element Communication Protocol (PCEP) Extensions for Network Resource Partition (NRP)", Work in Progress, Internet-Draft, draft-dong-pce-pcep-nrp-01, 23 October 2023, <https://www.ietf.org/archive/id/draft-dong-pce-pcep-nrp-01.txt>.

[I-D.ietf-idr-sr-policy-nrp] Dong, J., Hu, Z., and R. Pang, "BGP SR Policy Extensions for Network Resource Partition", Work in Progress,Internet-Draft, draft-ietf-idr-sr-policy-nrp-00, 17 December 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-nrp-00>.

[I-D.draft-ietf-idr-bgp-ls-sr-policy-nrp] R. Chen, J. Dong, D. Zhao, L. Gong, Y. Zhu and R. Pang, "SR Policies Extensions for Network Resource Partition in BGP-LS", Work in Progress,Internet-Draft, draft-ietf-idr-bgp-ls-sr-policy-nrp-00,02 October 2024,< https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ls-sr-policy-nrp-00>.

Acknowledgements

The authors would like to thank XXX for the review and discussion of this document.

Authors' Addresses

Shengnan Yue
China Mobile
China
yueshengnan@chinamobile.com

Ran Chen
ZTE Corporation
Email: chen.ran@zte.com.cn

Jie Dong
Huawei Technologies
Beijing
China
Email: jie.dong@huawei.com

Changwang Lin
New H3C Technologies
Beijing
China
Email: linchangwang.04414@h3c.com

Ran Pang
China Unicom
Beijing
China
Email: pangran@chinaunicom.cn

Ka Zhang
Huawei Technologies
Beijing
China
Email: zhangka@huawei.com

Wei Gao
CAICT
Beijing
China
Email: gaowei@caict.ac.cn

Wenying Jiang
China Mobile
Beijing
China

Email: jiangwenying@chinamobile.com