

IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: 16 April 2026

J. Dong
Z. Hu
Huawei Technologies
R. Pang
China Unicom
13 October 2025

BGP SR Policy Extensions for Network Resource Partition
draft-ietf-idr-sr-policy-nrp-04

Abstract

Segment Routing (SR) Policy is a set of candidate paths, each consisting of one or more segment lists and the associated information. The header of a packet steered in an SR Policy is augmented with an ordered list of segments associated with that SR Policy. A Network Resource Partition (NRP) is a subset of network resources allocated in the underlay network which can be used to support one or a group of RFC 9543 network slice services.

In networks where there are multiple NRPs, an SR Policy may be associated with a particular NRP. The association between SR Policy and NRP needs to be specified, so that for service traffic which is steered into the SR Policy, the header of the packets can be augmented with the information associated with the NRP. An SR Policy candidate path can be distributed using BGP SR Policy. This document defines the extensions to BGP SR policy to specify the NRP which the SR Policy candidate path is associated with.

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 16 April 2026.

Copyright Notice

Copyright (c) 2025 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
1.1. Requirements Language	3
2. NRP Identifier of SR Policy	3
3. Procedures	5
4. Error Handling	6
5. Scalability Considerations	6
6. Security Considerations	6
7. IANA Considerations	7
8. Acknowledgments	7
9. References	7
9.1. Normative References	7
9.2. Informative References	8
Authors' Addresses	9

1. Introduction

The concept of Segment Routing (SR) policy is defined in [RFC9256]. An SR Policy is a set of candidate paths, each consisting of one or more segment lists. The head end of an SR Policy may learn multiple candidate paths for an SR Policy. The header of a packet steered in an SR Policy is augmented with an ordered list of segments associated with that SR Policy. The BGP extensions to distribute SR Policy candidate paths is defined in [RFC9830].

[RFC9543] discusses the general framework, components, and interfaces for requesting and operating network slices using IETF technologies. It also introduces the concept Network Resource Partition (NRP), which is a subset of the resources and associated policies in the underlay network. The network slices defined in [RFC9543] can be realized by mapping one or more connectivity constructs to an NRP. [RFC9732] describes the framework and the candidate component technologies for providing enhanced VPN services based on VPN and

Traffic Engineering (TE) technologies. Enhanced VPN (VPN+) can be used for the realization of network slice services defined in [RFC9543].

As described in [I-D.ietf-teas-nrp-scalability], one scalable data plane approach to support network slicing is to carry a dedicated NRP 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 networks where there are multiple NRPs, an SR Policy may be associated with a particular NRP. The association between SR Policy and NRP needs to be specified, so that for service traffic which is steered into the SR Policy, the header of the packets can be augmented with the information associated with the NRP. The association between SR Policy and NRP is described in [I-D.jiang-spring-sr-policy-nrp]. An SR Policy candidate path can be distributed using BGP SR Policy. This document defines the extensions to BGP SR policy to specify the NRP which the SR Policy candidate path is associated with.

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 BCP14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. NRP Identifier of SR Policy

In order to specify the NRP the candidate path of SR policy is associated with, a new sub-TLV called "NRP" sub-TLV is defined in the BGP Tunnel Encapsulation Attribute [RFC9012]. The NRP sub-TLV can be carried in the BGP Tunnel Encapsulation Attribute with the tunnel type set to SR Policy. The use of the NRP sub-TLV in other tunnel types is outside the scope of this document.

The NRP sub-TLV has the following format:




Figure 1. NRP Sub-TLV

where:

- * Type: 123 (assigned by IANA)
- * Length: 6 octets.
- * Flags: 1-octet flag field. None is defined at this stage. The flags SHOULD be set to zero on transmission and MUST be ignored on receipt.
- * Reserved: 1 octet of reserved bits. It SHOULD be set to zero on transmission and MUST be ignored on receipt.
- * NRP ID: A 32-bit domain significant identifier which is used to identify an NRP in the control plane. The values of 0 and 0xFFFFFFFF are reserved.

The validation of an SR Policy NLRI with the NRP Sub-TLV in the BGP tunnel encapsulation attribute [RFC9012] follows the procedures in section 4.2 of [RFC9830], augmented by the validation procedures described in this document.

When the NRP sub-TLV is carried in the BGP Tunnel Encapsulation Attribute of an SR Policy NLRI, a segment list of the candidate path is considered invalid if the headend node of the SR Policy determines that the set of network resources corresponding to the NRP ID on network segments identified by the segment list do not exist. The detailed mechanisms for NRP resource validation are out of the scope of this document.

The encoding structure of BGP SR Policy with the NRP sub-TLV is expressed as below:

```

SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>
Attributes:
  Tunnel Encaps Attribute (23)
    Tunnel Type: SR Policy (15)
      Binding SID
      SRv6 Binding SID
      Preference
      Priority
      Policy Name
      Policy Candidate Path Name
      Explicit NULL Label Policy (ENLP)
      NRP
      Segment List
        Weight
        Segment
        Segment
        ...
      ...

```

Figure 2. SR Policy Encoding with NRP sub-TLV

3. Procedures

When a candidate path of SR Policy is instantiated within an NRP, and a network-wide data plane NRP Selector ID is used for identifying the resources of the NRP, the originating node of SR Policy SHOULD include the NRP sub-TLV in the BGP Tunnel Encapsulation Attribute of the BGP SR Policy. The setting of other fields and attributes in BGP SR Policy SHOULD follow the mechanism as defined in [RFC9830].

On reception of an SR Policy NLRI, a BGP speaker determines if it is acceptable and usable according to the rules defined in Section 4.2 of [RFC9830] and section 2 of this document. 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 map the NRP ID to the data plane NRP Selector ID, then encapsulate both the NRP Selector ID and the segment list of the selected candidate path in the header of packets which are steered to the SR Policy. For SR Policy with IPv6 data plane, the data plane NRP Selector ID can be the same as the NRP ID, and the approach to encapsulate the NRP Selector ID in IPv6 Hop-by-Hop Options header is defined in [I-D.ietf-6man-enhanced-vpn-vtn-id]. For SR Policy with MPLS data plane, the mechanisms of mapping and encapsulation of the NRP Selector ID in the packet would be based on the framework defined in [RFC9789].

Although the proposed mechanism allows that different candidate paths in one SR policy be associated with different NRPs, in normal network scenarios it is considered that the association between an SR Policy and NRP is consistent, in such case all candidate paths of one SR policy SHOULD be associated with the same NRP.

4. Error Handling

The error handling of the BGP Update messages for BGP SR Policy SAFI with the NRP extensions defined in this document follows the procedures in section 5 of [RFC9830].

The NRP sub-TLV is optional and MUST NOT appear more than once for one SR Policy candidate path. The NRP sub-TLV is considered malformed if its format does not match the above description. If the NRP sub-TLV appears more than once, or its format is considered malformed, the associated BGP SR Policy NLRI is considered malformed and the "treat-as-withdraw" strategy of [RFC7606] MUST be applied.

5. Scalability Considerations

The mechanism specified in this document adds additional information to the SR Policy candidate paths. In order to steer traffic into different NRPs using SR Policy, the SR Policies used for different NRPs need to be different. As the number of NRP increases, the number of SR Policies would also increase accordingly. When BGP is used for distributing SR Policy candidate paths, the amount of control plane information exchanged between the network controller and the headend nodes would also increase. However, since the SR Policies candidate paths distributed in BGP are only installed by the corresponding headend nodes, the impacts to the BGP control plane are considered acceptable.

6. Security Considerations

The security considerations of BGP [RFC4271] and BGP SR policy [RFC9830] apply to this document.

The NRP sub-TLV provides the NRP identifier that may be carried in IPv6 Hop-by-Hop options header or used in the encapsulation of MPLS. This NRP identifier can impact packet forwarding in a network so care should be taken to protect this mission-critical or commercially sensitive information during provisioning, query and report of the NRP-ID in BGP.

7. IANA Considerations

IANA has assigned the sub-TLV type as defined in Section 2 from "BGP Tunnel Encapsulation Attribute sub-TLVs" registry.

Value	Description	Reference
123	NRP	This document

8. Acknowledgments

The authors would like to thank Guoqi Xu, Lei Bao, Haibo Wang, Shunwan Zhuang and Susan Hares for their review and discussion of this document.

9. References

9.1. 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>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", RFC 7606, DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/info/rfc7606>>.
- [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>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", 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>>.
- [RFC9732] Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for NRP-Based Enhanced Virtual Private Networks", RFC 9732, DOI 10.17487/RFC9732, March 2025, <<https://www.rfc-editor.org/info/rfc9732>>.
- [RFC9830] Previdi, S., Filsfils, C., Talaulikar, K., Ed., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", RFC 9830, DOI 10.17487/RFC9830, September 2025, <<https://www.rfc-editor.org/info/rfc9830>>.

9.2. Informative References

- [I-D.ietf-6man-enhanced-vpn-vtn-id]
Dong, J., Li, Z., Xie, C., Ma, C., and G. S. Mishra, "Carrying Network Resource (NR) related Information in IPv6 Extension Header", Work in Progress, Internet-Draft, draft-ietf-6man-enhanced-vpn-vtn-id-12, 7 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-6man-enhanced-vpn-vtn-id-12>>.
- [I-D.ietf-teas-nrp-scalability]
Dong, J., Li, Z., Gong, L., Yang, G., and G. S. Mishra, "Scalability Considerations for Network Resource Partition", Work in Progress, Internet-Draft, draft-ietf-teas-nrp-scalability-07, 2 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-nrp-scalability-07>>.
- [I-D.jiang-spring-sr-policy-nrp]
Yue, Chen, R., Dong, J., Lin, C., Pang, R., KaZhang, gao, W., and J. Wenying, "Segment Routing Policy Extension for Network Resource Partition", Work in Progress, Internet-Draft, draft-jiang-spring-sr-policy-nrp-03, 30 June 2025, <<https://datatracker.ietf.org/doc/html/draft-jiang-spring-sr-policy-nrp-03>>.
- [RFC9789] Andersson, L., Bryant, S., Bocci, M., and T. Li, "MPLS Network Actions (MNAs) Framework", RFC 9789, DOI 10.17487/RFC9789, July 2025, <<https://www.rfc-editor.org/info/rfc9789>>.

Authors' Addresses

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

Zhibo Hu
Huawei Technologies
Email: huzhibo@huawei.com

Ran Pang
China Unicom
Email: pangran@chinaunicom.cn