Workgroup: IDR Working Group
Internet-Draft:
draft-dong-idr-sr-policy-nrp-02
Published: 31 January 2023
Intended Status: Standards Track
Expires: 4 August 2023

Authors: J. Dong — Huawei Technologies
Z. Hu — Huawei Technologies
R. Pang — China Unicom

BGP SR Policy Extensions for Network Resource Partition

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 IETF 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

Copyright Notice

Table of Contents

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 [I-D.ietf-idr-segment-routing-te-policy].

[I-D.ietf-teas-ietf-network-slices] introduces the concept and the characteristics of IETF network slice, and describes a general framework for IETF network slice management and operation. It also introduces the concept Network Resource Partition (NRP), which is a subset of the resources and associated policies in the underlay network. IETF network slice can be realized by mapping one or more connectivity constructs to an NRP. [I-D.ietf-teas-enhanced-vpn] describes the framework and the candidate component technologies for providing enhanced VPN (VPN+) services based on VPN and Traffic

Engineering (TE) technologies. Enhanced VPN (VPN+) can be used for the realization of IETF network slices. In the context of network slicing, an NRP is considered as an instantiation of the VTN as defined in [I-D.ietf-teas-enhanced-vpn].

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. 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 NRP sub-TLV is optional and MUST NOT appear more than once for one SR Policy candidate path. If the NRP sub-TLV appears more than once, the associated BGP SR Policy NLRI is considered malformed and the "treat-as-withdraw" strategy of [RFC7606] is applied.

The NRP sub-TLV has the following format:

```
   0                   1                   2                   3
   0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
  +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
  |     Type      |   Length      |     Flags     |   Reserved    |
  +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
  |                         NRP ID (4 octets)                     |
  +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 1. NRP Sub-TLV

where:

*Type: 123

*Length: 6

*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. Value 0 and 0xFFFFFFFF are reserved.

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 with a specific 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 [I-D.ietf-idr-segment-routing-te-policy].

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 [I-D.ietf-idr-segment-routing-te-policy]. 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 the NRP 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 approach to encapsulate the NRP 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, one approach to encapsulate the NRP ID to the packet is defined in [I-D.li-mpls-enhanced-vpn-vtn-id].

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. Security Considerations

The security considerations of BGP [RFC4271] and BGP SR policy [I-D.ietf-idr-segment-routing-te-policy] apply to this document.

5. 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

6. Acknowledgments

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

7. References

7.1. Normative 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://www.ietf.org/archive/id/draft-ietf-idr-segment-routing-te-policy-20.txt>.

[I-D.ietf-teas-enhanced-vpn] Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Enhanced Virtual Private Network (VPN+)", Work in Progress, Internet-Draft, draft-ietf-teas-enhanced-vpn-12, 23 January 2023, <https://www.ietf.org/archive/id/draft-ietf-teas-enhanced-vpn-12.txt>.

[I-D.ietf-teas-ietf-network-slices]
Farrel, A., Drake, J., Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura, "A Framework for IETF Network Slices", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slices-19, 21 January 2023, <https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slices-19.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>.

[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>.

7.2. Informative References

[I-D.ietf-6man-enhanced-vpn-vtn-id] Dong, J., Li, Z., Xie, C., Ma, C., and G. S. Mishra, "Carrying Virtual Transport Network (VTN) Information in IPv6 Extension Header", Work in Progress, Internet-Draft, draft-ietf-6man-enhanced-vpn-vtn-id-02, 24 October 2022, <https://www.ietf.org/archive/id/draft-ietf-6man-enhanced-vpn-vtn-id-02.txt>.

[I-D.ietf-teas-nrp-scalability]
Dong, J., Li, Z., Gong, L., Yang, G., Guichard, J., Mishra, G. S., Qin, F., Saad, T., and V. P. Beeram, "Scalability Considerations for Network Resource Partition", Work in Progress, Internet-Draft, draft-ietf-teas-nrp-scalability-01, 24 October 2022, <https://www.ietf.org/archive/id/draft-ietf-teas-nrp-scalability-01.txt>.

[I-D.li-mpls-enhanced-vpn-vtn-id] Li, Z. and J. Dong, "Carrying Virtual Transport Network (VTN) Information in MPLS Packet", Work in Progress, Internet-Draft, draft-li-mpls-enhanced-vpn-vtn-id-03, 16 October 2022, <https://www.ietf.org/archive/id/draft-li-mpls-enhanced-vpn-vtn-id-03.txt>.

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