

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 14 September 2023

Z. Li
J. Dong
Huawei Technologies
R. Pang
China Unicom
Y. Zhu
China Telecom
13 March 2023

Segment Routing for End-to-End IETF Network Slicing
draft-li-spring-sr-e2e-ietf-network-slicing-06

Abstract

IETF network slices can be used to meet the connectivity and performance requirements of different services or customers in a shared network. An IETF network slice can be realized by mapping a set of connectivity constructs to a network resource partition (NRP). In some network scenarios, an end-to-end IETF network slice may span multiple network domains. Within each domain, traffic of the end-to-end network slice service is mapped to an intra-domain NRP.

When segment routing (SR) is used to provide multi-domain IETF network slices, information of the intra-domain NRP can be specified using special SR binding segments which are called NRP binding segments (NRP BSID). Then a multi-domain IETF network slice can be specified using a list of NRP BSIDs in the packet, each of which is used by the corresponding domain edge nodes to steer the traffic of the end-to-end IETF network slice into the specific intra-domain NRP.

This document describes the functionality of the NRP binding segment and its instantiation in SR-MPLS and SRv6 data planes.

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

Copyright (c) 2023 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. Segment Routing for IETF E2E Network Slicing | 4 |
| 3. SRv6 NRP Behaviors | 5 |
| 3.1. End.B6NRP.Encaps | 5 |
| 3.2. End.NRP.Encaps | 6 |
| 3.3. End.BNRP.Encaps | 7 |
| 4. SR-MPLS NRP BSIDs | 9 |
| 5. IANA Considerations | 10 |
| 6. Security Considerations | 11 |
| 7. Acknowledgements | 11 |
| 8. References | 11 |
| 8.1. Normative References | 11 |
| 8.2. Informative References | 12 |
| Authors' Addresses | 13 |

1. Introduction

[I-D.ietf-teas-ietf-network-slices] introduces the concept and the characteristics of IETF network slices, and describes a general framework for IETF network slice management and operation. It also introduces the concept of the Network Resource Partition (NRP), which is a collection of resources identified in the underlay network. An IETF network slice can be realized by mapping a set of connectivity constructs to a network resource partition (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.

[I-D.ietf-teas-nrp-scalability] describes the scalability considerations in the control plane and data plane of NRPs and provides suggestions to improve the scalability of NRPs. In the data plane, it proposes to carry an NRP-ID in the data packet to determine the set of resources reserved for the corresponding NRP. [I-D.ietf-6man-enhanced-vpn-vtn-id] describes the mechanism of carrying the VTN resource ID (which is equivalent to NRP-ID) of a network domain in the IPv6 Hop-by-Hop (HBH) extension header.

An end-to-end IETF network slice may span multiple network domains. Within each domain, traffic of the end-to-end network slice service needs to be mapped to an intra-domain NRP. On the domain edge nodes, the NRP in the local domain used for carrying the end-to-end network slice needs to be determined. [I-D.li-teas-composite-network-slices] introduces the network slice related identifiers with different network scope. It also describes the approach of mapping the multi-domain NRP-ID to the intra-domain NRP-IDs at the network domain border nodes.

In SR networks, an NRP can be established and represented using either a set of NRP-specific resource-aware segments [I-D.ietf-spring-resource-aware-segments] [I-D.ietf-spring-sr-for-enhanced-vpn], or an NRP-ID which can identify the set of network resources allocated to an NRP.

When segment routing (SR) is used to provide end-to-end IETF network slices, information of the intra-domain NRP can be specified using special SR binding segments called NRP binding segments and indicated by Segment Identifiers (SIDs) called NRP binding segment identifiers (NRP BSID). Then an inter-domain NRP can be specified using a list of NRP BSIDs in the packet, each of which is used by the corresponding domain edge nodes to steer the traffic of the end-to-end IETF network slice into the specific intra-domain NRP.

This document describes the functionality of the NRP binding segment and its instantiation in SR-MPLS and SRv6 data plane.

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. Segment Routing for IETF E2E Network Slicing

With Segment Routing, there are several optional approaches to steer the end-to-end network slice traffic into the intra-domain NRPs. These approaches can be classified into two categories.

The first category of the approaches are to use an NRP BSID to steer traffic to an SR Policy which is associated with an intra-domain NRP. This is called the NRP Traffic Engineering (NRP-TE) BSID. There are two variants in terms of the detailed behavior:

- * The first variant is to use an NRP BSID to specify the mapping of traffic to an SR policy which consists of list of resource-aware segments [I-D.ietf-spring-resource-aware-segments] associated with a intra-domain NRP.
- * The second variant is to use an NRP BSID to specify the mapping of traffic to an SR policy which is associated with an intra-domain NRP-ID.

The second category of approaches are to use an NRP BSID to steer traffic to follow the shortest path within an intra-domain NRP. This is called the NRP Best Effort (NRP-BE) BSID. There are two variants in terms of the detailed behavior:

- * The first variant is to use an NRP BSID to determine an intra-domain NRP-ID, and instruct the domain edge node to encapsulate the intra-domain NRP-ID into the packet.
- * The second variant is to use an NRP BSID to specify the mapping of traffic to an intra-domain NRP, the intra-domain NRP-ID is specified in some fields of the packet by the ingress node of the end-to-end network slice, and is obtained and encapsulated into the packet at the domain edge node.

The behavior of the NRP-TE BSID is similar to the function of the existing SR BSID, the difference is that it is associated with a particular intra-domain NRP. The behavior of the NRP-BE BSID is different from the existing SR BSID. The instantiation of the NRP BSIDs in SR-MPLS and SRv6 are described in the following sections.

3. SRv6 NRP Behaviors

[RFC8986] defines the SRv6 Network Programming concept and specifies the base set of SRv6 behaviors. The SRv6 End.B6.Encaps behavior is defined to bind to an SRv6 Policy with encapsulation, and it can be used for the first variant of the NRP-TE BSID. In this case, the SRv6 End.B6 encaps function is used to steer the network slice traffic to an SRv6 Policy, which consists of candidate paths built with resource-aware SRv6 segment lists that are associated with an intra-domain NRP.

For other types and variants of NRP binding segments as described in section 2, three new SRv6 behaviors are defined as shown in the following subsections.

3.1. End.B6NRP.Encaps

A new SRv6 function called End.B6NRP.Encaps: Endpoint bound to an SRv6 Policy with IPv6 NRP encapsulation is defined. This is a variation of the End behavior. It instructs the endpoint node to determine an SRv6 Policy in a specific NRP of the local-domain, and encapsulate both the SID list and the NRP-ID specified by the SRv6 Policy in a new IPv6 header.

Any SID instance of this behavior is associated with an SR Policy B, an NRP-ID V, and a source address A.

When node N receives a packet whose IPv6 DA is S, and S is a local End.B6NRP.Encaps SID, N does the following:

```
S01. When an SRH is processed {
S02.   If (Segments Left == 0) {
S03.     Stop processing the SRH, and proceed to process the next
           header in the packet, whose type is identified by
           the Next Header field in the routing header.
S04.   }
S05.   If (IPv6 Hop Limit <= 1) {
S06.     Send an ICMP Time Exceeded message to the Source Address
           with Code 0 (Hop limit exceeded in transit),
           interrupt packet processing, and discard the packet.
S07.   }
S08.   max_LE = (Hdr Ext Len / 2) - 1
S09.   If ((Last Entry > max_LE) or (Segments Left > Last Entry+1)) {
S10.     Send an ICMP Parameter Problem to the Source Address
           with Code 0 (Erroneous header field encountered)
           and Pointer set to the Segments Left field,
           interrupt packet processing, and discard the packet.
S11.   }
S12.   Decrement IPv6 Hop Limit by 1
S13.   Decrement Segments Left by 1
S14.   Update IPv6 DA with Segment List[Segments Left]
S15.   Push a new IPv6 header with its own SRH containing B, and
           set the NRP-ID in the HBH header to V
S16.   Set the outer IPv6 SA to A
S17.   Set the outer IPv6 DA to the first SID of B
S18.   Set the outer Payload Length, Traffic Class, Flow Label,
           Hop Limit, and Next Header fields
S19.   Submit the packet to the egress IPv6 FIB lookup for
           transmission to the new destination
S20. }
```

Note:

Comparing with the End.B6.Encaps behavior, the difference is in step 15, which includes the setting of the NRP-ID in the IPv6 HBH header

3.2. End.NRP.Encaps

A new SRv6 function called End.NRP.Encaps: Endpoint with IPv6 NRP encapsulation is defined. This is a variation of the End behavior. It instructs the endpoint node to determine the corresponding NRP-ID of the local domain based on the mapping relationship between the End.NRP.Encaps SID and the intra-domain NRPs maintained on the endpoint. Then the NRP-ID is carried in the IPv6 HBH header of the packet.

Any SID instance of this behavior is associated with an NRP-ID V.

When node N receives a packet whose IPv6 DA is S, and S is a local End.NRP.Encaps SID, N does the following:

```
S01. When an SRH is processed {
S02.   If (Segments Left == 0) {
S03.     Stop processing the SRH, and proceed to process the next
           header in the packet, whose type is identified by
           the Next Header field in the routing header.
S04.   }
S05.   If (IPv6 Hop Limit <= 1) {
S06.     Send an ICMP Time Exceeded message to the Source Address
           with Code 0 (Hop limit exceeded in transit),
           interrupt packet processing, and discard the packet.
S07.   }
S08.   max_LE = (Hdr Ext Len / 2) - 1
S09.   If ((Last Entry > max_LE) or (Segments Left > Last Entry+1)) {
S10.     Send an ICMP Parameter Problem to the Source Address
           with Code 0 (Erroneous header field encountered)
           and Pointer set to the Segments Left field,
           interrupt packet processing, and discard the packet.
S11.   }
S12.   Decrement IPv6 Hop Limit by 1
S13.   Decrement Segments Left by 1
S14.   Update IPv6 DA with Segment List [Segments Left]
S15.   Set the NRP-ID in the HBH header to V
S16.   Submit the packet to the egress IPv6 FIB lookup for
           transmission to the new destination
S17. }
```

Note:

Comparing with the End.B6NRP.Encaps behavior, the difference is in step 15 to 17, which does not need to include an SRH in the IPv6 header

3.3. End.BNRP.Encaps

A new SRv6 function called End.BNRP.Encaps: Endpoint bound to an IPv6 NRP with encapsulation is defined. This is a variation of the End behavior. For the End.BNRP SID, its corresponding NRP-ID is specified by the ingress node of the SRv6 path of the inter-domain NRP, and is carried in some fields of the packet. It instructs the endpoint node to obtain the corresponding NRP-ID from the received packet, and encapsulate it into the IPv6 HBH header of the packet for further forwarding. Through the End.BNRP.Encaps behavior, the ingress node can flexibly specify the intra-domain NRPs the packet needs to traverse in the multi-domain network.

Any SID instance of this behavior is associated with an NRP-ID V.

There can be several options to carry the intra-domain NRP-ID corresponding to the End.BNRP.Encaps behavior:

1. The NRP-ID is carried in the argument field of the End.BNRP.Encaps SID.
2. The NRP-ID is carried in the SRH TLV field.
3. The NRP-ID is carried in the next SID following the End.BNRP.Encaps SID in the SID list.

Editor's note: In the current version of this document, the option 1 is further specified. The use of other options is for further study.

When an ingress node of an end-to-end SR path of the inter-domain NRP encapsulates an End.BNRP.Encaps SID in the SID list, it SHOULD put the intra-domain NRP-ID which the packet is expected to be steered to in that domain into the argument part of the corresponding SID.

Any SID instance of this behavior contains one NRP-ID V in its argument.

When node N receives a packet whose IPv6 DA is S, and S is a local End.BNRP.Encaps SID, N does the following:


```
S01. When an SRH is processed {
S02.   If (Segments Left == 0) {
S03.     Stop processing the SRH, and proceed to process the next
           header in the packet, whose type is identified by
           the Next Header field in the routing header.
S04.   }
S05.   If (IPv6 Hop Limit <= 1) {
S06.     Send an ICMP Time Exceeded message to the Source Address
           with Code 0 (Hop limit exceeded in transit),
           interrupt packet processing, and discard the packet.
S07.   }
S08.   max_LE = (Hdr Ext Len / 2) - 1
S09.   If ((Last Entry > max_LE) or (Segments Left > Last Entry+1)) {
S10.     Send an ICMP Parameter Problem to the Source Address
           with Code 0 (Erroneous header field encountered)
           and Pointer set to the Segments Left field,
           interrupt packet processing, and discard the packet.
S11.   }
S12.   Obtain the NRP-ID V from the argument part of the IPv6 DA
S13.   Decrement IPv6 Hop Limit by 1
S14.   Decrement Segments Left by 1
S15.   Update IPv6 DA with Segment List [Segments Left]
S16.   Set the NRP-ID in the HBH header to V
S17.   Submit the packet to the egress IPv6 FIB lookup for
           transmission to the new destination
S18. }
```

Note:

Comparing with the End.NRP.Encaps behavior, the difference is in the new step 12, which is to obtain the NRP-ID from the current IPv6 DA.

4. SR-MPLS NRP BSIDs

[I-D.li-mpls-enhanced-vpn-vtn-id] describes the mechanism of carrying the VTN ID in the MPLS extension header. The VTN ID is equivalent to an NRP-ID.

With the SR-MPLS data plane, SR-MPLS BSIDs can be allocated by a domain edge node for different NRP Binding behaviors described in section 2.

For the first variant of NRP-TE BSID, an SR-MPLS BSID is bound to an SR Policy which consists of candidate paths built with resource-aware segment lists associated with an intra-domain NRP. When a node receives a packet with a locally assigned NRP-TE BSID, it determines the corresponding segment list which consists of the resource-aware segments of a intra-domain NRP, and encapsulates the SID list to the MPLS label stack.

For the second variant of the NRP-TE BSID, an SR-MPLS BSID is bound to an SR Policy associated with an intra-domain NRP-ID. When a node receives a packet with a locally assigned NRP-TE BSID, it determines the corresponding SID list and the intra-domain NRP-ID, and encapsulates the packet with both the SID list and an MPLS VTN extension header which carries the intra-domain NRP-ID. Note this requires to assign a separate NRP BSID for each SR policy in the intra-domain NRPs which the node participates in.

For the first variant of the NRP-BE BSID, an SR-MPLS BSID is bound to the shortest path in an intra-domain NRP. When a node receives a packet with a locally assigned NRP-BE BSID, it determines the corresponding intra-domain NRP-ID based on the mapping relationship between the NRP-BE BSID and the intra-domain NRPs, and encapsulates the packet with an MPLS VTN extension header which carries the intra-domain NRP-ID. Note this requires to assign a separate NRP-BE BSID for each intra-domain NRP.

For the second variant of the NRP-BE BSID, an SR-MPLS BSID is bound to the shortest path in an intra-domain NRP, the NRP-ID is specified by the E2E SR path ingress node of the inter-domain NRP and is carried in the MPLS VTN extension header. When a node receives a packet with a locally assigned NRP-BE BSID, it obtains the corresponding intra-domain NRP-ID from an NRP-ID list carried in the packet, then encapsulates the obtained intra-domain NRP-ID into the MPLS VTN extension header of the packet.

5. IANA Considerations

IANA is requested to assign the following code points from the "SRv6 Endpoint Behaviors" sub-registry in the "Segment-routing with IPv6 data plane (SRv6) Parameters" registry:

| Value | Hex | Endpoint Behavior | Reference |
|-------|-----|-------------------|-----------|
| TBA1 | | End.BNRP.Encaps | [This ID] |
| TBA2 | | End.NRP | [This ID] |
| TBA3 | | End.BNRP | [This ID] |

6. Security Considerations

The security considerations of segment routing [RFC8402] [RFC8754] applies to this document.

7. Acknowledgements

The authors would like to thank Zhibo Hu and Yawei Zhang for their review and valuable comments.

8. References

8.1. Normative References

[I-D.ietf-spring-resource-aware-segments]

Dong, J., Bryant, S., Miyasaka, T., Zhu, Y., Qin, F., Li, Z., and F. Clad, "Introducing Resource Awareness to SR Segments", Work in Progress, Internet-Draft, draft-ietf-spring-resource-aware-segments-06, 11 October 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-resource-aware-segments-06>>.

[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://datatracker.ietf.org/doc/html/draft-ietf-teas-enhanced-vpn-12>>.

[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://datatracker.ietf.org/doc/html/draft-ietf-teas-ietf-network-slices-19>>.

[I-D.li-teas-composite-network-slices]

Li, Z., Dong, J., Pang, R., Zhu, Y., and L. M. Contreras, "Realization of Composite IETF Network Slices", Work in Progress, Internet-Draft, draft-li-teas-composite-network-slices-00, 13 March 2023, <<https://datatracker.ietf.org/api/v1/doc/document/draft-li-teas-composite-network-slices/>>.

- [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>>.
- [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>>.

8.2. 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://datatracker.ietf.org/doc/html/draft-ietf-6man-enhanced-vpn-vtn-id-02>>.
- [I-D.ietf-spring-sr-for-enhanced-vpn]
Dong, J., Bryant, S., Miyasaka, T., Zhu, Y., Qin, F., Li, Z., and F. Clad, "Segment Routing based Virtual Transport Network (VTN) for Enhanced VPN", Work in Progress, Internet-Draft, draft-ietf-spring-sr-for-enhanced-vpn-04, 11 October 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-sr-for-enhanced-vpn-04>>.
- [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://datatracker.ietf.org/doc/html/draft-ietf-teas-nrp-scalability-01>>.
- [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://datatracker.ietf.org/doc/html/draft-li-mpls-enhanced-vpn-vtn-id-03>>.

- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.

Authors' Addresses

Zhenbin Li
Huawei Technologies
Huawei Campus, No. 156 Beiqing Road
Beijing
100095
China
Email: lizhenbin@huawei.com

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Road
Beijing
100095
China
Email: jie.dong@huawei.com

Ran Pang
China Unicom
Email: pangran@chinaunicom.cn

Yongqing Zhu
China Telecom
Email: zhuyq8@chinatelecom.cn