

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 25 September 2022

Z. Li
J. Dong
Huawei Technologies
R. Pang
China Unicom
Y. Zhu
China Telecom
24 March 2022

Segment Routing for End-to-End IETF Network Slicing
draft-li-spring-sr-e2e-ietf-network-slicing-03

Abstract

Network slicing can be used to meet the connectivity and performance requirement of different services or customers in a shared network. An IETF network slice can be realized as enhanced VPNs (VPN+), which is delivered by integrating the overlay VPN service with a Virtual Transport Network (VTN) as the underlay. 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 a domain VTN. In the context of IETF network slicing, a VTN can be instantiated as a Network Resource Partition (NRP).

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

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

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 [RFC 2119](#) [[RFC2119](#)].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Draft

SR for E2E IETF Network Slicing

March 2022

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 25 September 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	3
2.	Segment Routing for IETF E2E Network Slicing	4
3.	SRv6 NRP Binding Functions	5
3.1.	End.B6NRP.Encaps	5
3.2.	End.NRP.Encaps	6
3.3.	End.BNRP.Encaps	7
4.	SR-MPLS NRP BSIDs	8
5.	IANA Considerations	9
6.	Security Considerations	9
7.	Acknowledgements	9
8.	References	10
8.1.	Normative References	10
8.2.	Informative References	10
	Authors' Addresses	11

1. Introduction

[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 collection of resources identified in the underlay network.

[I-D.ietf-teas-enhanced-vpn] describes the framework and the candidate component technologies for providing enhanced VPN (VPN+) services based on existing VPN and Traffic Engineering (TE) technologies with enhanced characteristics that specific services require above traditional VPNs. It also introduces the concept of Virtual Transport Network (VTN). A Virtual Transport Network (VTN) is a virtual underlay network which consists of a set of dedicated or shared network resources allocated from the physical underlay network, and is associated with a customized logical network topology. VPN+ services can be delivered by mapping one or a group of overlay VPNs to the appropriate VTNs as the underlay, so as to provide the network characteristics required by the customers. Enhanced VPN (VPN+) and VTN can be used for the realization of IETF network slices. In the context of IETF network slicing, a VTN can be instantiated as an NRP. VTN and NRP are considered interchangeable terms in this document.

[I-D.dong-teas-nrp-scalability] describes the scalability considerations in the control plane and data plane to enable NRPs and provide the suggestions to improve the scalability of NRP. In the control plane, It proposes the approach of decoupling the topology and resource attributes of NRP, so that multiple NRPs may share the same topology and the result of topology based path computation. In the data plane, it proposes to carry a dedicated NRP-ID of a network domain in the data packet to determine the set of resources reserved for the corresponding NRP.

An IETF network slice may span multiple network domains. Within each

domain, traffic of the end-to-end network slice is mapped to a local network slice. The NRP ID which identifies the NRP in the local domain for the end-to-end network slice needs to be determined on the domain edge node.

When segment routing (SR) is used to build a multi-domain IETF network slice, information of the local network slices in each domain can be specified using special SR binding segments called NRP binding segments (NRP BSID). The multi-domain IETF network slice can be specified using a list of NRP BSIDs in the packet, each of which can be used by the corresponding domain edge nodes to steer the traffic of end-to-end IETF network slice using the specific resource-aware segments or NRP-ID of the local domain.

This document describes the functionality of the network slice binding segment and its instantiation in SR-MPLS and SRv6.

[2.](#) Segment Routing for IETF E2E Network Slicing

[I-D.dong-teas-nrp-scalability] describes the scalability considerations in the control plane and data plane to create NRPs. In data plane, it proposes to carry a dedicated NRP-ID in data packet to determine the set of resources reserved for the corresponding NRP in a network domain.

[I-D.li-teas-e2e-ietf-network-slicing] describes the framework of carrying network slice related identifiers in the data plane, each of the network slice IDs may have a different network scope. It provides an approach of mapping the global NRP-ID to domain NRP-IDs at the network domain border nodes.

With Segment Routing, there are several optional approaches to realize the mapping between the end-to-end network slice and the network slice constructs in the local domain.

The first type of approaches are to use one type of NRP BSID to steer traffic to an SR Policy associated with a local NRP. This is called the NRP-TE BSID. There are some variants in terms of the detailed behavior:

- * The first variant is to use one type of NRP BSID to specify the mapping of traffic to a SR policy which consists of list of resource-aware segments [[I-D.ietf-spring-resource-aware-segments](#)] associated with a local NRP.
- * The second variant is to use one type of NRP BSID to specify the mapping of traffic to a SR policy which is bound to a local NRP-ID.

The second type of approaches is to use one type of NRP BSID to steer traffic to follow the shortest path within a local domain NRP. This is called the NRP-BE BSID. There are some variants in terms of the detailed behavior:

- * The first variant is to use one type of NRP BSID to determine a local NRP-ID, and instruct the encapsulation of the local NRP-ID into the packet at the domain edge node.
- * The second variant is to use one type of NRP BSID to specify the mapping of traffic to a local NRP, the local NRP-ID is specified in the associated fields by the ingress node, and is encapsulated into the packet at the domain edge node.

The behavior of the first type of NRP BSID is similar to the function of the existing SR BSID, the difference is it is associated with a particular NRP. The second type of the NRP BSID is different from the existing binding segment. The instantiation of the NRP BSIDs in SR-MPLS and SRv6 are described in the following sections.

[3.](#) SRv6 NRP Binding Functions

[RFC8986] defines the SRv6 Network Programming concept and specifies the base set of SRv6 behaviors. The SRv6 End.B6.Encaps function is defined to instantiate the Binding SID in SRv6, which can be reused as one type of NRP-TE BSID to specify the mapping of traffic to a list of resource-aware SRv6 segments of a domain NRP.

[I-D.ietf-6man-enhanced-vpn-vtn-id] describes the mechanism of carrying the VTN-ID of a network domain in the IPv6 Hop-by-Hop (HBH) extension header. For the type 2, 3, 4 of NRP binding segments described in [section 2](#), three new SRv6 Binding functions are defined in the following sections.

[3.1.](#) End.B6NRP.Encaps

A new SRv6 function called End.B6NRP.Encaps: Endpoint bound to a SRv6 Policy in a NRP with IPv6 encapsulation is defined in this section. 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 the SID list of the SR Policy and the NRP-ID in a new IPv6 header.

Any SID instance of this behavior is associated with an SR Policy B, a 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
      the VTN-ID in VTN option set to V in the HBH Ext header
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. }
```

[3.2.](#) End.NRP.Encaps

A new SRv6 function called End.NRP.Encaps 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 NRPs maintained on the endpoint. The NRP-ID is encapsulated in the VTN option in the IPv6 HBH extension header.

Any SID instance of this behavior is associated with one 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.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 VTN-ID in VTN option to V in the HBH Ext header
S16.    Submit the packet to the egress IPv6 FIB lookup for
            transmission to the new destination
S17. }

```

[3.3.](#) End.BNRP.Encaps

A new SRv6 function called End.BNRP.Encaps: Endpoint bound to a NRP with IPv6 encapsulation is defined. This is a variation of the End behavior. For the End.BNRP SID, its corresponding NRP-ID should be specified and encapsulated by the ingress node of SRv6 Path. It instructs the endpoint node to obtain the corresponding NRP-ID from the SRH, and encapsulate it in the VTN option in the IPv6 HBH extension header. Through the End.BNRP.Encaps, the ingress node can flexibly specify the local NRP the packet traverses in the network.

Any SID instance of this behavior is associated with one NRP-ID V and a source address A.

There can be several options to carry the local NRP-ID corresponding to the End.BNRP.Encaps function:

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, option 1 is preferred, in which the local NRP-ID is carried in the argument field of the SRv6 SID.

When an ingress node of an SR path encapsulates the End.BNRP.Encaps SID into the packet, it SHOULD put the NRP-ID which the packet is expected to be mapped to into the argument part of the SID.

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 VTN-ID in VTN option to V in the HBH Ext header
S17.   Submit the packet to the egress IPv6 FIB lookup for
           transmission to the new destination
S18. }
```

[4.](#) SR-MPLS NRP BSIDs

[I-D.li-mpls-enhanced-vpn-vtn-id] describes the mechanism of carrying the VTN-ID of a network domain in the MPLS extension header.

With SR-MPLS data plane, NRP BSIDs can be allocated by a domain edge node for the three types of NRP binding behaviors described in [section 2](#).

For the first type of NRP BSID, a BSID can be bound to a list of resource-aware segments of a local NRP. When a node receives a packet with a locally assigned NRP BSID, it determines the corresponding SID list which consists of the resource-aware segments of a local NRP, and encapsulates the SID list to the MPLS label stack.

For another variant of the first type NRP BSID, a NRP BSID is bound to a SR Policy and a local NRP-ID. When a node receives a packet with a locally assigned NRP BSID, it determines the corresponding SID list and the local NRP-ID, and encaps the packet with the SID list and an MPLS VTN extension header which carries the local NRP-ID. Note this requires to assign a NRP BSID for each SR policy in each NRP the node participates in.

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

For a variant of the second type NRP BSID, a NRP BSID is bound to the shortest path in an NRP of the local network domain, the NRP-ID is specified and encapsulated by the ingress node in the MPLS VTN extension header. When a node receives a packet with a locally assigned NRP BSID, it obtains the corresponding local NRP-ID from the NRP-ID list in the VTN extension header, and update the local NRP-ID in the VTN extension header with the obtained NRP-ID.

[5.](#) IANA Considerations

TBD

[6.](#) Security Considerations

TBD

[7.](#) Acknowledgements

The authors would like to thank Zhibo Hu for his review and valuable

[8.](#) References

[8.1.](#) Normative References

[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+) Services", Work in Progress, Internet-Draft, [draft-ietf-teas-enhanced-vpn-10](#), 6 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-teas-enhanced-vpn-10.txt>>.

[I-D.ietf-teas-ietf-network-slices]

Farrel, A., Drake, J., Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura, "Framework for IETF Network Slices", Work in Progress, Internet-Draft, [draft-ietf-teas-ietf-network-slices-08](#), 6 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slices-08.txt>>.

[I-D.li-teas-e2e-ietf-network-slicing]

Li, Z., Dong, J., Pang, R., and Y. Zhu, "Framework for End-to-End IETF Network Slicing", Work in Progress, Internet-Draft, [draft-li-teas-e2e-ietf-network-slicing-02](#), 7 March 2022, <<https://www.ietf.org/archive/id/draft-li-teas-e2e-ietf-network-slicing-02.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>>.

[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.dong-teas-nrp-scalability]

Dong, J., Li, Z., Gong, L., Yang, G., Guichard, J. N., Mishra, G., Qin, F., Saad, T., and V. P. Beeram, "Scalability Considerations for Network Resource Partition", Work in Progress, Internet-Draft, [draft-dong-teas-nrp-scalability-01](https://www.ietf.org/archive/id/draft-dong-teas-nrp-scalability-01), 7 February 2022, <<https://www.ietf.org/archive/id/draft-dong-teas-nrp-scalability-01.txt>>.

Li, et al.

Expires 25 September 2022

[Page 10]

Internet-Draft

SR for E2E IETF Network Slicing

March 2022

[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](https://www.ietf.org/archive/id/draft-ietf-6man-enhanced-vpn-vtn-id-00), 5 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-6man-enhanced-vpn-vtn-id-00.txt>>.

[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-04](https://www.ietf.org/archive/id/draft-ietf-spring-resource-aware-segments-04), 5 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-spring-resource-aware-segments-04.txt>>.

[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-02](https://www.ietf.org/archive/id/draft-ietf-spring-sr-for-enhanced-vpn-02), 5 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-spring-sr-for-enhanced-vpn-02.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](https://www.ietf.org/archive/id/draft-li-mpls-enhanced-vpn-vtn-id-02), 7 March 2022, <<https://www.ietf.org/archive/id/draft-li-mpls-enhanced-vpn-vtn-id-02.txt>>.

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

Li, et al.

Expires 25 September 2022

[Page 11]

Internet-Draft

SR for E2E IETF Network Slicing

March 2022

Email: jie.dong@huawei.com

Ran Pang
China Unicom
Email: pangran@chinaunicom.cn

Yongqing Zhu
China Telecom
Email: zhuyq8@chinatelecom.cn

