

Workgroup: Network Working Group

Internet-Draft:

draft-li-spring-sr-e2e-ietf-network-slicing-01

Published: 25 October 2021

Intended Status: Standards Track

Expires: 28 April 2022

Authors: Z. Li

J. Dong

Huawei Technologies

Huawei Technologies

R. Pang

Y. Zhu

China Unicom

China Telecom

Segment Routing for End-to-End IETF Network Slicing

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 local VTN.

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 VTN binding segments (VTN BSID). The multi-domain IETF network slice can be specified using a list of VTN 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 VTN in the local domain.

This document describes the functionality of VTN 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-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 28 April 2022.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- 1. [Introduction](#)
- 2. [Segment Routing for IETF E2E Network Slicing](#)
- 3. [SRv6 VTN Binding Functions](#)
 - 3.1. [End.VTN.Encaps](#)
 - 3.2. [End.BVTN.Encaps](#)
 - 3.3. [End.B6VTN.Encaps](#)
- 4. [SR-MPLS VTN BSIDs](#)
- 5. [IANA Considerations](#)
- 6. [Security Considerations](#)
- 7. [Acknowledgements](#)
- 8. [References](#)
 - 8.1. [Normative References](#)
 - 8.2. [Informative References](#)
- [Authors' Addresses](#)

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.

[[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.

[[I-D.dong-teas-enhanced-vpn-vtn-scalability](#)] describes the scalability considerations in the control plane and data plane to enable VPN+ services, and provide the suggestions to improve the scalability of VTN. In the control plane, It proposes the approach of decoupling the topology and resource attributes of VTN, so that multiple VTNs may share the same topology and the result of topology based path computation. In the data plane, it proposes to carry a VTN-ID of a network domain in the data packet to determine the set of resources reserved for the corresponding VTN.

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 VTN-ID which identifies the virtual underlay network 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 VTN binding segments (VTN BSID). The multi-domain IETF network slice can be specified using a list of VTN 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 VTN-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-enhanced-vpn-vtn-scalability](#)] describes the scalability considerations in the control plane and data plane to enable VPN+ services. In data plane, it proposes to carry a dedicated VTN-ID in data packet to determine the set of resources reserved for the corresponding VTN 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 VTN-ID to local VTN-IDs at the network domain edge 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.

1. The first approach is to use one type of VTN binding segment to specify the mapping of traffic to a list of resource-aware segments of a local VTN.
2. The second approach is to use one type of VTN binding segment to determine a local VTN-ID, and instruct the encapsulation of the local VTN-ID into the packet at the domain edge node.
3. The third approach is to use one type of VTN binding segment to specify the mapping of traffic to a local VTN, the local VTN-ID is specified in the associated fields by the ingress node, and is encapsulated into the packet at the domain edge node.
4. The fourth approach is to use one type of VTN binding segment to specify the mapping of traffic to a SR policy which is bound to a local VTN.

The function of the first type of VTN binding segment is similar to the function of the existing binding segment, the difference is it is associated with a particular VTN. The other types of the VTN binding segments are different from the existing binding segment, and their instantiation in SR-MPLS and SRv6 are described in the following sections.

3. SRv6 VTN 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 used as the first type of VTN binding function to specify the mapping of traffic to a list of resource-aware SRv6 segments of a local VTN.

[[I-D.dong-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 VTN binding segments described in section 2, three new SRv6 Binding functions are defined in the following sections.

3.1. End.VTN.Encaps

A new SRv6 function called End.VTN.Encaps is defined. This is a variation of the End behavior. It instructs the endpoint node to determine the corresponding VTN-ID of the local domain based on the

mapping relationship between the End.VTN.Encaps SID and the VTNs maintained on the endpoint. The VTN-ID is encapsulated in the VTN-ID option in the IPv6 HBH extension header.

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

When node N receives a packet whose IPv6 DA is S, and S is a local End.VTN.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 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. }
```

3.2. End.BVTN.Encaps

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

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

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

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

Editor's note: In the current version of this document, option 1 is preferred, in which the local VTN-ID is carried in the argument field of the SRv6 SID.

When an ingress node of an SR path encapsulates the End.BVTN.Encaps SID into the packet, it SHOULD put the VTN-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.BVTN.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 VTN-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 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. }
```

3.3. End.B6VTN.Encaps

A new SRv6 function called End.B6VTN.Encaps: Endpoint bound to a SRv6 Policy in a VTN 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 VTN of the local domain, and encapsulate the SID list of the SR Policy and the VTN-ID in a new IPv6 header.

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

When node N receives a packet whose IPv6 DA is S, and S is a local End.B6VTN.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 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. }
```

4. SR-MPLS VTN 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, VTN binding SIDs can be allocated by a domain edge node for the three types of VTN binding behaviors described in section 2.

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

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

For the third type of VTN binding segment, a VTN BSID is bound to a VTN of the local network domain, the VTN-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 VTN BSID, it obtains the corresponding local VTN-ID from the VTN-ID list in the VTN extension header, and update the local VTN-ID option in the VTN extension header with the obtained VTN-ID.

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

5. IANA Considerations

TBD

6. Security Considerations

TBD

7. Acknowledgements

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

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-08, 12 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-teas-enhanced-vpn-08.txt>>.

[I-D.ietf-teas-ietf-network-slices] Farrel, A., Gray, E., 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-04, 23 August 2021, <<https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slices-04.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-01, 25 October 2021, <<https://www.ietf.org/archive/id/draft-li-teas-e2e-ietf-network-slicing-01.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-6man-enhanced-vpn-vtn-id] Dong, J., Li, Z., Xie, C., Ma, C., and G. Mishra, "Carrying Virtual Transport Network Identifier in IPv6 Extension Header", Work in Progress, Internet-Draft, draft-dong-6man-enhanced-vpn-vtn-id-05, 8 September 2021, <<https://www.ietf.org/archive/id/draft-dong-6man-enhanced-vpn-vtn-id-05.txt>>.

[I-D.dong-teas-enhanced-vpn-vtn-scalability]

Dong, J., Li, Z., Gong, L., Yang, G., Guichard, J. N., Mishra, G., and F. Qin, "Scalability Considerations for Enhanced VPN (VPN+)", Work in Progress, Internet-Draft, draft-dong-teas-enhanced-vpn-vtn-scalability-03, 11 July 2021, <<https://www.ietf.org/archive/id/draft-dong-teas-enhanced-vpn-vtn-scalability-03.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-03, 12 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-spring-resource-aware-segments-03.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-01, 12 July 2021, <<https://www.ietf.org/archive/id/draft-ietf-spring-sr-for-enhanced-vpn-01.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-01, 14 April 2021, <<https://www.ietf.org/archive/id/draft-li-mpls-enhanced-vpn-vtn-id-01.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

Email: jie.dong@huawei.com

Ran Pang
China Unicom

Email: pangran@chinaunicom.cn

Yongqing Zhu
China Telecom

Email: zhuyq8@chinatelecom.cn