Workgroup: Network Working Group
Internet-Draft:
draft-li-mpls-enhanced-vpn-vtn-id-03
Published: 16 October 2022
Intended Status: Standards Track
Expires: 19 April 2023
Authors: Z. Li
Huawei Technologies
J. Dong
Huawei Technologies

Carrying Virtual Transport Network (VTN) Information in MPLS Packet

Abstract

A Virtual Transport Network (VTN) is a virtual network which has a customized network topology and a set of dedicated or shared network resources allocated from the underlying physical network. Multiple VTNs can be created by network operator for using as the underlay for one or a group of VPNs services to provide enhanced VPN (VPN+) services. In packet forwarding, some fields in the data packet needs to be used to identify the VTN the packet belongs to, so that the VTN-specific processing can be executed on the packet. In the context of network slicing, a VTN can be instantiated as a Network Resource Partition (NRP).

This document proposes a mechanism to carry the data plane VTN ID in an MPLS packet to identify the VTN the packet belongs to. The procedure for processing the VTN ID is also specified.

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 19 April 2023.

Table of Contents

1. Introduction

Virtual Private Networks (VPNs) provide different groups of users with logically isolated connectivity over a common shared network infrastructure. With the evolution of 5G and cloud, new service types may require connectivity services with advanced characteristics comparing to traditional VPNs, such as strict isolation from other services or guaranteed performance. These services are referred to as "enhanced VPNs" (VPN+). [I-D.ietf-teas-enhanced-vpn] describes a framework and candidate component technologies for delivering VPN+ services.

The enhanced properties of VPN+ require integration between the overlay connectivity and the resource and characteristics provided by the underlay network. To meet the requirement of VPN+ services, a number of Virtual Transport Networks (VTNs) need to be created, each has a logical network topology and a set of network resources allocated from the underlay network to meet the requirement of one or a group of VPN+ services. In the network, traffic of different VPN+ services may to be processed separately according to the topology and the network resources associated with the corresponding

VTN. [I-D.ietf-teas-ietf-network-slices] introduces the concept Network Resource Partition (NRP) as a set of network resources that are available to carry traffic and meet the SLOs and SLEs. In the context of network slicing, a VTN can be instantiated as a Network Resource Partition (NRP).

For network scenarios where a large number of VTNs need to be created and maintained, [I-D.ietf-teas-nrp-scalability] describes the scalability considerations for NRP. One approach to improve the data plane scalability is introducing a dedicated VTN Identifier (VTN ID) in data packets to identify the VTN the packets belong to, so that VTN-specific packet processing can be performed by network nodes.

This document proposes a mechanism to carry the VTN Identifier (VTN ID) and the related information in MPLS [RFC3031] data packets, so that the packet will be processed by network nodes using the set of network resources allocated to the corresponding VTN. The procedure for processing the VTN ID is also specified. The destination and forwarding path of the MPLS LSP is determined using the MPLS label stack in the packet, and the set of local network resources used for processing the packet is determined by the VTN ID. The mechanism introduced in this document is applicable to both MPLS networks with RSVP-TE [RFC3209] or LDP [RFC5036] LSPs, and MPLS networks with Segment Routing (SR) [RFC8402] [RFC8660].

2. 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.

3. Carrying VTN Information in MPLS Packet

This document makes use of the post stack extension header mechanism as defined in [I-D.song-mpls-extension-header]. A new type of MPLS extension header called "VTN extension header" is defined to carry the VTN ID and other VTN related information. The type of VTN extension header is to be assigned by IANA. The format of VTN extension header is shown as below:

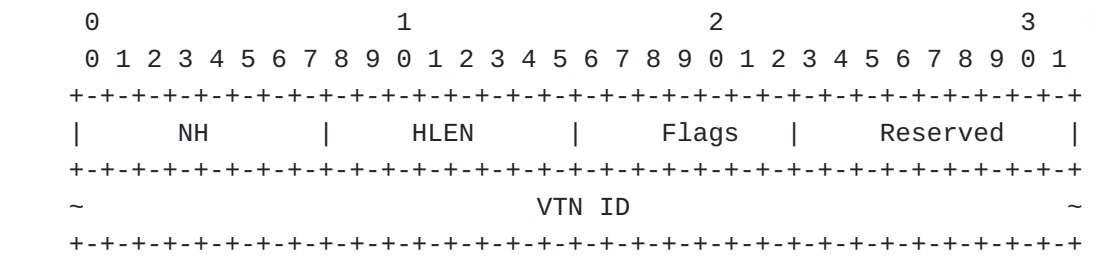

Figure 1. The format of MPLS VTN Extension Header

Where:

*NH: 8-bit indicator for the Next Header type.

*HLEN: 8-bit unsigned integer for the Extension Header Length in 4-octet units, not including the first 4 octets.

*Flags: 8-bit flag field. None of them is allocated in this version.

*Reserved: 8-bit field reserved for future use.

*VTN ID: 4-octet identifier which uniquely identifies the set of network resources allocated to a VTN.

The VTN extension header SHOULD be processed hop-by-hop (HBH). Thus it is suggested the VTN extension header be positioned in precedence over the end-to-end (E2E) extension headers.

The benefit of introducing the post-stack VTN Extension Header to carry the VTN ID and related information is that it provides the flexibility to encode information which cannot be accommodated in an MPLS label stack, and the length of the extension header can be variable.

4. Procedures

4.1. VTN Extension Header Insertion

When the ingress node of an LSP receives a packet, according to traffic classification or mapping policy, the packet is steered into one of the VTNs in the network, then an MPLS VTN extension header SHOULD be inserted into the Post-Stack extension headers of the packet, and the VTN ID which the packet is mapped to SHOULD be carried in the VTN header. The ingress node SHOULD also encapsulates the packet with an MPLS label stack which are used to determine the destination and path of the LSP.

4.2. VTN based Packet Forwarding

On receipt of a MPLS packet which carries the VTN extension header, network nodes which support the mechanism defined in this document SHOULD parse the VTN header and use the VTN ID to identify the VTN the packet belongs to, and use the local resources allocated to the VTN to process and forward the packet. The forwarding behavior is based on both the MPLS label stack and the VTN extension header. The top MPLS label is used for the lookup of the next-hop, and the VTN ID can be used to determine the set of network resources allocated by the network nodes for processing and sending the packet to the next-hop.

There can be different approaches used for allocating network resources on each network node to the VTNs. For example, on one interface, a subset of forwarding plane resource (e.g., bandwidth and the associated buffer/queuing/scheduling resources) allocated to a particular VTN can be considered as a virtual layer-2 sub-interface with dedicated bandwidth and the associated resources. In packet forwarding, the top MPLS label of the received packet is used to identify the next-hop and the outgoing Layer 3 interface, and the VTN-ID is used to further identify the virtual sub-interface which is associated with the VTN on the outgoing interface.

Network nodes which do not support the mechanism in this document SHOULD ignore the VTN extension header, and forward the packet only based on the top MPLS label.

The egress node of the MPLS LSP SHOULD pop the VTN extension header, together with other post-stack extension headers if there is any.

5. Capability Advertisement and Negotiation

Before inserting the VTN extension header into an MPLS packet, the ingress node MAY want to know whether the nodes along the LSP can process the VTN extension header properly based on the mechanisms defined in this document. This can be achieved by introducing the capability advertisement and negotiation mechanism for the VTN extension header. The ingress node also need to know whether the egress node of the LSP can remove the VTN extension header as part of the post-stack action header properly before parsing the upper layer and send the packet to the next hop. The capability advertisement and negotiation mechanism will be described in a future version of this document.

6. IANA Considerations

IANA is requested to assign the code point for a new type of extension header as below:

Value	Description	Reference
TBD	VTN	this document

7. Security Considerations

TBD

8. Contributors

Zhibo Hu
Email: huzhibo@huawei.com

9. Acknowledgements

TBD.

10. References

10.1. 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+)", Work in Progress, Internet-Draft, draft-ietf-teas-enhanced-vpn-11, 19 September 2022, <https://www.ietf.org/archive/id/draft-ietf-teas-enhanced-vpn-11.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-14, 3 August 2022, <https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slices-14.txt>.

[I-D.song-mpls-extension-header] Song, H., Zhou, T., Andersson, L., Zhang, Z., and R. Gandhi, "MPLS Network Actions using Post-Stack Extension Headers", Work in Progress, Internet-Draft, draft-song-mpls-extension-header-11, 15 October 2022, <https://datatracker.ietf.org/api/v1/doc/document/draft-song-mpls-extension-header/>.

[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>.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/

RFC3031, January 2001, <https://www.rfc-editor.org/info/rfc3031>.

[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>.

10.2. Informative References

[I-D.ietf-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-ietf-teas-nrp-scalability-00, 11 July 2022, <https://www.ietf.org/archive/id/draft-ietf-teas-nrp-scalability-00.txt>.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <https://www.rfc-editor.org/info/rfc3209>.

[RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, DOI 10.17487/RFC5036, October 2007, <https://www.rfc-editor.org/info/rfc5036>.

[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>.

[RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019, <https://www.rfc-editor.org/info/rfc8660>.

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