

Carrying Virtual Transport Network Identifier in MPLS Packet
draft-li-mpls-enhanced-vpn-vtn-id-01

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 network infrastructure. 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.

This document proposes a mechanism to carry the 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 October 16, 2021.

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
2.	Requirements Language	3
3.	Carrying VTN Information in MPLS Packet	3
4.	Procedures	5
4.1.	VTN Header Insertion	5
4.2.	VTN based Packet Forwarding	5
5.	Capability Advertisement and Negotiation	6
6.	IANA Considerations	6
7.	Security Considerations	6
8.	Contributors	6
9.	Acknowledgements	6
10.	References	7
10.1.	Normative References	7
10.2.	Informative References	7
	Authors' Addresses	8

[1.](#) Introduction

Virtual Private Networks (VPNs) provide different groups of users with logically isolated connectivity over a common shared network infrastructure. With the introduction of 5G, 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 providing VPN+ services.

The enhanced properties of VPN+ require integration between the overlay connectivity and the characteristics provided by the underlay network. To meet the requirement of enhanced VPN services, a number of Virtual Transport Networks (VTNs) need to be created, each consists of a subset of the underlay 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 based on the topology and the network resources associated with the corresponding VTN.

For network scenarios where a large number of VTNs need to be created and maintained, [[I-D.dong-teas-enhanced-vpn-vtn-scalability](#)] describes the scalability considerations for VTN. 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 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 RFC 2119](#) [[RFC2119](#)] [RFC 8174](#) [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

3. Carrying VTN Information in MPLS Packet

This document defines a new VTN header which is used to carry the VTN-ID and other VTN related information. In an MPLS packet, The VTN header follows the MPLS label stack, and precedes the header and payloads in the upper layer. The format of VTN header is shown as below:




Figure 1. The format of MPLS VTN Header

Where:

- o Nibble: The first 4-bit field is set to the binary value 0010. This is to ensure that the VTN header will not be interpreted as an IP header or the ACH of pseudowire packet.
- o Length: Indicate the length of the MPLS VTN header in 32-bit words.
- o Flags: 8-bit Flags field. All the flags are reversed for future use. This field SHOULD be set to zero on transmission and MUST be ignored on receipt.
- o Reserved: 16-bit field reserved for future use.

A new VTN-ID Option is defined in this document, other option types may be defined in future documents. The format of the VTN-ID Option is shown as below:




Figure 2. The format of VTN-ID Option

Option Type: 8-bit identifier of the type of option. The type of VTN-ID option is to be assigned by IANA. The highest-order bits of the type field are defined as below:

- o BB 00 The highest-order 2 bits are set to 00 to indicate that a node which does not recognize this type will skip over it and continue processing the header.
- o C 1 The third highest-order bit are set to 1 to indicate this option may change en route.

Opt Data Len: 8-bit unsigned integer indicates the length of the option Data field of this option, in octets. The value of Opt Data Len of the VTN-ID option SHOULD be set to 4.

Option Data: 4-octet identifier which uniquely identifies a VTN within a network domain.

A new MPLS special-purpose label or extended special-purpose label is defined as the VTN Header Indicator (VHI), its value is to be assigned by IANA. The VHI label is used to indicate the existence of the VTN Header after the MPLS label stack in the packet. The position of the VHI label in the MPLS label stack is not limited.

The benefit of introducing the MPLS VTN header to carry the VTN-ID and the related information is that it provides the flexibility to encode information which cannot be accommodated in an MPLS label (20-bit), and the length of the header can be variable.

4. Procedures

4.1. VTN 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 a VTN header SHOULD be inserted into 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 path traversed by the LSP. The VHI label SHOULD be inserted in the label stack to identify the existence of the VTH header.

4.2. VTN based Packet Forwarding

On receipt of a MPLS packet which carries the VHL and the VTN header, network nodes which support the mechanism defined in this document SHOULD scan the label stack to figure out the existence of the VHL. If there is a VHL in the label stack, then the network node 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 top MPLS label and the VTN-ID. 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 VHL and the VTN header, and forward the packet only based on the top MPLS label.

The egress node of the MPLS LSP SHOULD pop the VHL together with other LSP labels, and decapsulate the VTN header.

5. Capability Advertisement and Negotiation

Before inserting the VTN header into an MPLS packet, the ingress node MAY need to know whether the nodes along the LSP can process the VTN header properly according to the mechanisms defined in this document. This can be achieved by introducing the capability advertisement and negotiation mechanism for the VTN header. The ingress node also need to know whether the egress node of the LSP can remove the VTN 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 a new special-purpose label from the "Special-Purpose MPLS Label Values" or "Extended Special-Purpose MPLS Label Values" registry.

Value	Description	Reference

TBD	VTN Header Indicator	this document

IANA is requested to assign a new option type of the MPLS VTN extension header:

Value	Description	Reference

TBD	VTN-ID	this document

7. Security Considerations

TBD

8. Contributors

Zhibo Hu
Email: huzhibo@huawei.com

9. Acknowledgements

TBD.

10. References

10.1. Normative References

- [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>>.
- [RFC7274] Kompella, K., Andersson, L., and A. Farrel, "Allocating and Retiring Special-Purpose MPLS Labels", [RFC 7274](#), DOI 10.17487/RFC7274, June 2014, <<https://www.rfc-editor.org/info/rfc7274>>.
- [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.dong-teas-enhanced-vpn-vtn-scalability] Dong, J., Li, Z., Qin, F., and G. Yang, "Scalability Considerations for Enhanced VPN (VPN+)", [draft-dong-teas-enhanced-vpn-vtn-scalability-01](#) (work in progress), November 2020.
- [I-D.ietf-teas-enhanced-vpn] Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Enhanced Virtual Private Networks (VPN+) Service", [draft-ietf-teas-enhanced-vpn-06](#) (work in progress), July 2020.
- [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>>.
- [TS23501] "3GPP TS23.501", 2016, <<https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>>.

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

