Authors: J. Dong   Z. Li
         Huawei Technologies   Huawei Technologies
         C. Xie   C. Ma   G. Mishra
         China Telecom   China Telecom   Verizon Inc.

Carrying Virtual Transport Network (VTN) Information in IPv6 Extension Header

Abstract

Virtual Private Networks (VPNs) provide different customers with logically separated connectivity over a common network infrastructure. With the introduction and evolvement of 5G and other network scenarios, some existing or new customers may require connectivity services with advanced characteristics comparing to traditional VPNs. Such kind of network service is called enhanced VPNs (VPN+). VPN+ can be used to deliver IETF network slices, and could also be used for other application scenarios.

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 virtual underlay. 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 performed on each node the packet traverses.

This document proposes a new Hop-by-Hop option of IPv6 extension header to carry the VTN related information, which could used to identify the set of network resources allocated to a VTN and the rules for packet processing. The procedure for processing the VTN option 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 12 January 2023.

Table of Contents

1. **Introduction**

Virtual Private Networks (VPNs) provide different customers with
logically isolated connectivity over a common network
infrastructure. With the introduction and evolvement of 5G and other
network scenarios, some existing or new customers may require
connectivity services with advanced characteristics comparing to
traditional VPNs, such as resource isolation from other services or
guaranteed performance. Such kind of network service is called
enhanced VPN (VPN+). VPN+ service requires the coordination and

integration between the overlay VPNs and the capability and resources of the underlay network. VPN+ can be used to deliver IETF network slices [I-D.ietf-teas-ietf-network-slices].

[I-D.ietf-teas-enhanced-vpn] describes a framework and the candidate component technologies for providing VPN+ services. It also introduces the concept of Virtual Transport Network (VTN). A 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 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. In packet forwarding, traffic of different VPN+ services needs to be processed separately based on the network resources and the logical topology associated with the corresponding VTN. In the context of network slicing, VTN and NRP are considered as similar concepts, and NRP can be seen as an instantiation of VTN.

[I-D.dong-teas-nrp-scalability] describes the scalability considerations and the possible optimizations for providing a relatively large number of VTNs for VPN+ services. One approach to improve the data plane scalability of VTN is to introduce a dedicated VTN Resource Identifier (VTN Resource ID) in the data packet to identify the set of network resources allocated to a VTN, so that VTN-specific packet processing can be performed using that set of resources, which avoids the possible resource competition with services in other VTNs. This is called Resource Independent (RI) VTN. A VTN Resource ID represents a subset of the resources (e.g. bandwidth, buffer and queuing resources) allocated on a given set of links and nodes which constitute a logical network topology. The logical topology associated with a VTN could be defined using mechanisms such as Multi-Topology [RFC4915], [RFC5120] or Flex-Algo [I-D.ietf-lsr-flex-algo], etc.

This document proposes a mechanism to carry the VTN resource ID together with other VTN related information in a new Hop-by-Hop option called "VTN option" of IPv6 extension header [RFC8200] of IPv6 packet, so that on each network node along the packet forwarding path, the VTN option in the packet is parsed, and the obtained VTN Resource ID is used to instruct the network node to use the set of network resources allocated to the corresponding VTN to process and forward the packet. The procedure for processing the VTN option is also specified. This provides a scalable solution to support a relatively large number of VTNs in an IPv6 network.

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
BCP14 RFC 2119 [RFC2119] RFC 8174 [RFC8174] when, and only when,
they appear in all capitals, as shown here.

2. New IPv6 Extension Header Option for VTN

A new Hop-by-Hop option type "VTN" is defined to carry the VTN
resource identifier and other VTN related information in an IPv6
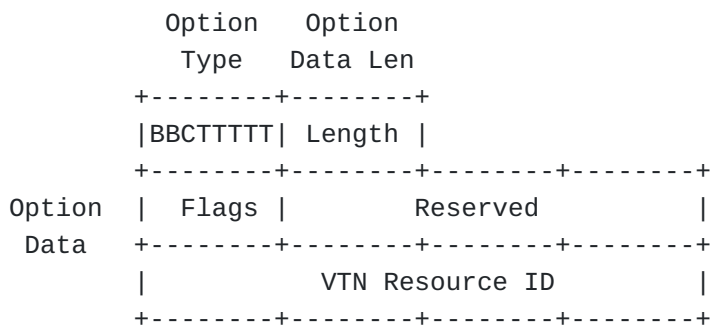packet. Its format is shown as below:

```
                   Option   Option
                    Type   Data Len
                  +--------+--------+
                  |BBCTTTTT| Length |
                  +--------+--------+--------+--------+
        Option   | Flags  |       Reserved          |
         Data    +--------+--------+--------+--------+
                 |            VTN Resource ID        |
                 +--------+--------+--------+--------+
            Figure 1. The format of VTN Option
```

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

  *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.

  *C 0 The third highest-order bit is set to 0 to indicate this
   option does not change en route.

  *TTTTT To be assigned by IANA.

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 VTN option SHOULD be set to 8.

Flags: 8-bit flags field. The most significant bit is defined in
this document.

```
       0 1 2 3 4 5 6 7
      +-+-+-+-+-+-+-+-+
      |S|U U U U U U U|
      +-+-+-+-+-+-+-+-+
```

   *S (Strict Match): The S flag is used to indicate whether the VTN
    Resource ID MUST be strictly matched for the processing of the
    packet. When S flag is set to 1, if the VTN resource ID in the
    VTN option does not match with any of the VTN resource ID
    provisioned on the outgoing interface of the network node, the
    packet MUST be dropped. When S flag is set to 0, if the VTN
    resource ID in the VTN option does not match with any of the VTN
    resource ID provisioned on the outgoing interface of the network
    node, the packet SHOULD be forwarded using the default set of
    network resource on the outgoing interface.

   *U (Unused): These flags are reserved for future use. They SHOULD
    be set to 0 on transmission and MUST be ignored on receipt.

   Reserved: 3-octet field reserved for future use. SHOULD be set to 0
   on transmission and MUST be ignored on receipt.

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

   Note that, if a deployment found it useful, the four-octet VTN
   Resource ID field may be derived from the four-octet Single Network
   Slice Selection Assistance Information (S-NSSAI) defined in 3GPP
   [TS23501].

3. Procedures

   As the VTN option needs to be processed by each node along the path
   for VTN-specific forwarding, it MUST be carried in IPv6 Hop-by-Hop
   options header.

3.1. Adding VTN Option to Packet

   When an ingress node of an IPv6 domain receives a packet, according
   to the traffic classification or mapping policy, the packet is
   steered into one of the VTNs in the network, then the packet MUST be
   encapsulated in an outer IPv6 header, and the Resource ID of the VTN
   which the packet is mapped to MUST be carried in the VTN option of
   the Hop-by-Hop Options header associated with the outer IPv6 header.

3.2. VTN based Packet Forwarding

   On receipt of a packet with the VTN option, each network node which
   can process the VTN option in fast path MUST use the VTN Resource ID
   to determine the set of local network resources allocated to the VTN

for packet processing. The packet forwarding behavior is based on both the destination IP address and the VTN Resource ID. More specifically, the destination IP address is used to determine the next-hop and the outgoing interface, and VTN Resource ID is used to determine the set of network resources on the outgoing interface which are allocated to the VTN for processing and sending the packet. If the VTN Resource ID does not match with any of the VTN Resource ID provisioned on the outgoing interface, the S flag in the VTN option is used to determine whether the packet is dropped or forwarded using the default set of network resources of the outgoing interface. The Traffic Class field of the outer IPv6 header MAY be used to provide differentiated treatment for packets which belong to the same VTN. The egress node of the IPv6 domain MUST decapsulate the outer IPv6 header and the Hop-by-Hop Options header which includes the VTN option.

In the forwarding plane, there can be different approaches of partitioning the local network resources and allocating them to different VTNs. For example, on one physical interface, a subset of the forwarding plane resources (e.g. bandwidth and the associated buffer and queuing resources) can be allocated to a particular VTN and represented as a virtual sub-interface with reserved bandwidth resource. In packet forwarding, the IPv6 destination address of the received packet is used to identify the next-hop and the outgoing layer-3 interface, and the VTN Resource ID is used to further identify the virtual sub-interface on the outgoing interface which is associated with the VTN.

Network nodes which do not support the processing of Hop-by-Hop Options header SHOULD ignore the Hop-by-Hop options header and forward the packet only based on the destination IP address. Network nodes which support Hop-by-Hop Options header, but do not support the VTN option SHOULD ignore the VTN option and continue to forward the packet based on the destination IP address. The network node MAY process the rest of the Hop-by-Hop options in the Hop-by-Hop Options header.

4. **Operational Considerations**

As described in [RFC8200], network nodes may be configured to ignore the Hop-by-Hop Options header, drop packets containing a Hop-by-Hop Options header, or assign packets containing a Hop-by-Hop Options header to a slow processing path. [I-D.ietf-6man-hbh-processing] specifies the modified procedures for the processing of IPv6 Hop-by-Hop Options header. Operator needs to make sure that all the network nodes involved in a VTN can either process the Hop-by-Hop Options header in the fast path, or ignore the Hop-by-Hop Options header. Since a VTN is associated with a logical network topology, it is practical to ensure that all the network nodes involved in that

logical topology support the processing of the Hop-by-Hop Options
header and the VTN option. In other word, packets steered into a VTN
MUST NOT be dropped due to the existence of the Hop-by-Hop Options
header. It is RECOMMENDED to configure all the network nodes
involved in a VTN to process the Hop-by-Hop Options header and the
VTN option if there is a nob for this.

5. IANA Considerations

This document requests IANA to assign a new option type from
"Destination Options and Hop-by-Hop Options" registry.

```
Value           Description       Reference
--------------------------------------------
 TBA            VTN Option      this document
```

6. Security Considerations

The security considerations with IPv6 Hop-by-Hop Options header are
described in [RFC8200], [RFC7045], [RFC9098] and [I-D.ietf-6man-hbh-
processing]. This document introduces a new IPv6 Hop-by-Hop option
which is either processed in the fast path or ignored by network
nodes, thus it does not introduce additional security issues.

7. Contributors

Zhibo Hu
Email: huzhibo@huawei.com

Lei Bao
Email: baolei7@huawei.com

8. Acknowledgements

The authors would like to thank Juhua Xu, James Guichard, Joel
Halpern, Tom Petch, Aijun Wang, Zhenqiang Li, Tom Herbert, Adrian
Farrel and Eric Vyncke for their review and valuable comments.

9. References

9.1. Normative References

[I-D.ietf-6man-hbh-processing] Hinden, R. M. and G. Fairhurst, "IPv6
          Hop-by-Hop Options Processing Procedures", Work in
          Progress, Internet-Draft, draft-ietf-6man-hbh-
          processing-01, 7 July 2022, <https://www.ietf.org/
          archive/id/draft-ietf-6man-hbh-processing-01.txt>.

[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>.

[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>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6
           (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/
           RFC8200, July 2017, <https://www.rfc-editor.org/info/
           rfc8200>.

9.2. 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-02, 16 May 2022, <https://
           www.ietf.org/archive/id/draft-dong-teas-nrp-
           scalability-02.txt>.

[I-D.ietf-lsr-flex-algo] Psenak, P., Hegde, S., Filsfils, C.,
           Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm",
           Work in Progress, Internet-Draft, draft-ietf-lsr-flex-
           algo-20, 18 May 2022, <https://www.ietf.org/archive/id/
           draft-ietf-lsr-flex-algo-20.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-12, 30 June 2022,
           <https://www.ietf.org/archive/id/draft-ietf-teas-ietf-
           network-slices-12.txt>.

[RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P.
           Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF",

RFC 4915, DOI 10.17487/RFC4915, June 2007, <https://www.rfc-editor.org/info/rfc4915>.

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi
           Topology (MT) Routing in Intermediate System to
           Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/
           RFC5120, February 2008, <https://www.rfc-editor.org/info/
           rfc5120>.

[RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing
           of IPv6 Extension Headers", RFC 7045, DOI 10.17487/
           RFC7045, December 2013, <https://www.rfc-editor.org/info/
           rfc7045>.

[RFC9098] Gont, F., Hilliard, N., Doering, G., Kumari, W., Huston,
           G., and W. Liu, "Operational Implications of IPv6 Packets
           with Extension Headers", RFC 9098, DOI 10.17487/RFC9098,
           September 2021, <https://www.rfc-editor.org/info/
           rfc9098>.

[RFC9099] Vyncke, É., Chittimaneni, K., Kaeo, M., and E. Rey,
           "Operational Security Considerations for IPv6 Networks",
           RFC 9099, DOI 10.17487/RFC9099, August 2021, <https://www.rfc-editor.org/info/rfc9099>.

[TS23501] "3GPP TS23.501", 2016, <https://portal.3gpp.org/
           desktopmodules/Specifications/SpecificationDetails.aspx?
           specificationId=3144>.

Authors' Addresses

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Road
Beijing
100095
China

Email: jie.dong@huawei.com

Zhenbin Li
Huawei Technologies
Huawei Campus, No. 156 Beiqing Road
Beijing
100095
China

Email: lizhenbin@huawei.com

Chongfeng Xie

China Telecom
China Telecom Beijing Information Science & Technology, Beiqijia
Beijing
102209
China

Email: xiechf@chinatelecom.cn

Chenhao Ma
China Telecom
China Telecom Beijing Information Science & Technology, Beiqijia
Beijing
102209
China

Email: machh@chinatelecom.cn

Gyan Mishra
Verizon Inc.

Email: gyan.s.mishra@verizon.com