

6man
Internet-Draft
Intended status: Standards Track
Expires: December 13, 2018

R. Bonica
Juniper Networks
C. Lenart
Verizon
G. Presbury
Hughes Network Systems
June 11, 2018

The IPv6 Virtual Private Network (VPN) Context Information Option
draft-bonica-6man-vpn-dest-opt-00

Abstract

This document defines a new IPv6 Destination Option that can be used to encode Virtual Private Network (VPN) context information. It is applicable when VPN payload is transported over IPv6.

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 December 13, 2018.

Copyright Notice

Copyright (c) 2018 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	4
3.	VPN Context Information	4
4.	The VPN Context Information Option	5
5.	Security Considerations	6
6.	IANA Considerations	6
7.	Acknowledgements	6
8.	References	6
8.1.	Normative References	6
8.2.	Informative References	7
	Authors' Addresses	8

[1.](#) Introduction

Virtual Private Network (VPN) technologies allow network providers to emulate private networks with shared infrastructure. For example, assume that a red sites and blue sites connect to a provider network. The provider network facilitates communication among red sites and facilitates communication among blue sites. However, it prevents communication between red sites and blue sites.

The IETF has standardized many VPN technologies, including:

- o Layer 2 VPN (L2VPN) [[RFC6624](#)].
- o Layer 3 VPN (L3VPN) [[RFC4364](#)].
- o Virtual Private LAN Service (VPLS) [[RFC4761](#)][[RFC4762](#)].
- o Ethernet VPN (EVPN) [[RFC7432](#)].
- o Pseudowires [[RFC8077](#)].

The above-mentioned technologies include the following components:

- o Customer Edge (CE) devices.
- o Provider Edge (PE) devices.
- o Routing Instances.
- o VPN context information.

- o Transport tunnels.

CE devices participate in closed communities called VPNs. CEs that participate in one VPN can communicate with one another but cannot communicate with CEs that participate in another VPN.

CE devices connect to provider networks through PE devices. Each PE maintains one Routing Instance for each VPN that it supports. A Routing Instance is a VPN specific Forwarding Information Base (FIB). In EVPN, Routing Instances are called Ethernet Virtual Instances (EVI).

Assume that one CE sends a packet through a provider network to another CE. The packet enters the provider network through an ingress PE and leaves the provider network through an egress PE. The packet may traverse one or more intermediate nodes on route from PE to PE.

When the ingress PE receives the packet, it:

- o Identifies the Routing Instance that supports the originating CE's VPN.
- o Searches that Routing Instance for the packet's destination.

If the search fails, the ingress PE discards the packet. If the search succeeds, it yields the following:

- o VPN context information.
- o The egress PE's IP address.

The ingress PE prepends VPN context information and a transport header to the packet, in that order. It then forwards the packet through a transport tunnel to the egress PE.

The egress PE removes the transport header, if it has not already been removed by an upstream device. It then examines and removes the VPN context information. Finally, it uses the VPN context information to forward the packet to its destination (i.e., a directly connected CE).

In the above-mentioned VPN technologies, the ingress PE encodes VPN context information in a Multiprotocol Label Switching (MPLS) [[RFC3031](#)] label. Depending upon the transport tunnel type, the transport header can be:

- o A MPLS label or label stack.

- o An IPv4 [[RFC0791](#)] header.
- o An IPv6 [[RFC8200](#)] header.
- o A Generic Routing Encapsulation (GRE) [[RFC2784](#)] header encapsulated in IPv4 or IPv6.

If the outermost transport header is IPv6, it may be followed by a Segment Routing Header (SRH) [[I-D.ietf-6man-segment-routing-header](#)].

Some PE devices cannot process MPLS headers. While these devices have several alternatives to MPLS-based transport tunnels, they require an alternative to MPLS-based encoding of VPN context information. This document defines a new IPv6 Destination Option that can be used to encode VPN context information. It is applicable when VPN payload is transported over IPv6.

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 [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

3. VPN Context Information

VPN context information specifies a forwarding procedure to be executed by the egress PE. However, VPN context information values are not globally mapped to forwarding procedures. Each egress PE maps each forwarding procedure that it supports to a VPN context information value. Therefore, VPN context information values are locally scoped to the egress PE.

PE devices can acquire VPN Context Information:

- o From one another, using a distributed, control plane protocol (e.g., BGP [[RFC4271](#)] [[RFC4760](#)])
- o From a controller.

The mechanisms by which PE devices acquire VPN Context Information are beyond the scope of this document.

4. The VPN Context Information Option




Figure 1

Figure 1 depicts the VPN Context Information Option. The IPv6 Destination Options header MAY include the VPN Context Information option. The IPv6 Hop-by-hop header MUST NOT include the VPN Context Information option.

Option fields are as follows:

- o Option Type - VPN Context Information option. Value TBD by IANA. See Notes below.
- o Opt Data Len - Length of Option Data, measured in bytes.
- o VPN Context Information - Specifies a forwarding procedure to be executed by the egress PE.

The VPN Context Information Option MUST NOT appear multiple times in a single packet. If a node receives a packet that contains multiple instances of the VPN Context Information Option, it MUST discard the packet and send an ICMP Parameter Problem, Code 2, message to the packet's source.

NOTE 1: The highest-order two bits of the Option Type (i.e., the "act" bits) are 10. These bits specify the action taken by a destination node that does not recognize VPN Context Information option. The required action is to discard the packet and, regardless of whether or not the packet's Destination Address was a multicast address, send an ICMPv6 [[RFC4443](#)] Parameter Problem, Code 2, message to the packet's Source Address, pointing to the unrecognized Option Type.

NOTE 2: The third highest-order bit of the Option Type (i.e., the "chg" bit) is 0. This indicates that Option Data cannot be modified along the path between the packet's source and its destination.

5. Security Considerations

A VPN can be deployed:

- o In a walled-garden environment.
- o In an over-the-top environment.

In a walled-garden environment, all PE devices and all devices that connect PEs to one another reside in the same security domain. Therefore, there is no risk that a packet might be modified as it travels from PE to PE.

In an over-the-top environment, all PE devices reside in one security domain while devices that connect PEs to one another can reside in a different security domain. In that case, there is significant risk that a packet might be modified as it travels from PE to PE.

Therefore, the VPN Context Information option MUST be authenticated when used in over-the-top environments. In this scenario, an IPv6 Encapsulating Security Payload (ESP) [RFC4303] MUST precede the Destination Options header that carries the VPN Context Information option. The ESP integrity service MUST be enabled.

6. IANA Considerations

IANA is requested to allocate a codepoint from the Destination Options and Hop-by-hop Options registry (<https://www.iana.org/assignments/ipv6-parameters/ipv6-parameters.xhtml#ipv6-parameters-2>). This option is called "VPN Context Information". The "act" bits are 10 and the "chg" bit is 0.

7. Acknowledgements

Thanks to Adrian Farrel and John Leddy for their comments.

8. References

8.1. Normative References

- [RFC0791] Postel, J., "Internet Protocol", STD 5, [RFC 791](#), DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.
- [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>>.

- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", [RFC 4303](#), DOI 10.17487/RFC4303, December 2005, <<https://www.rfc-editor.org/info/rfc4303>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, [RFC 4443](#), DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.
- [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>>.

8.2. Informative References

- [I-D.ietf-6man-segment-routing-header] Previdi, S., Filsfils, C., Leddy, J., Matsushima, S., and d. daniel.voyer@bell.ca, "IPv6 Segment Routing Header (SRH)", [draft-ietf-6man-segment-routing-header-13](#) (work in progress), May 2018.
- [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", [RFC 2784](#), DOI 10.17487/RFC2784, March 2000, <<https://www.rfc-editor.org/info/rfc2784>>.
- [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>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.

- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.
- [RFC4761] Kompella, K., Ed. and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", [RFC 4761](#), DOI 10.17487/RFC4761, January 2007, <<https://www.rfc-editor.org/info/rfc4761>>.
- [RFC4762] Lasserre, M., Ed. and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", [RFC 4762](#), DOI 10.17487/RFC4762, January 2007, <<https://www.rfc-editor.org/info/rfc4762>>.
- [RFC6624] Kompella, K., Kothari, B., and R. Cherukuri, "Layer 2 Virtual Private Networks Using BGP for Auto-Discovery and Signaling", [RFC 6624](#), DOI 10.17487/RFC6624, May 2012, <<https://www.rfc-editor.org/info/rfc6624>>.
- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", [RFC 7432](#), DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.
- [RFC8077] Martini, L., Ed. and G. Heron, Ed., "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", STD 84, [RFC 8077](#), DOI 10.17487/RFC8077, February 2017, <<https://www.rfc-editor.org/info/rfc8077>>.

Authors' Addresses

Ron Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, Virginia 20171
USA

Email: rbonica@juniper.net

Chris Lenart
Verizon
22001 Loudoun County Parkway
Ashburn, Virginia 20147
USA

Email: chris.lenart@verizon.com

Greg Presbury
Hughes Network Systems
11717 Exploration Lane
Germantown, Maryland 20876
USA

Email: greg.presbury@hughes.com