6man
Internet-Draft
Intended status: Experimental
Expires: 15 November 2025

R. Bonica
Juniper Networks
X. Li
CERNET Center/Tsinghua University
A. Farrel
Old Dog Consulting
Y. Kamite
NTT Communications Corporation
L. Jalil
Verizon

14 May 2025

The IPv6 VPN Service Destination Option
draft-ietf-6man-vpn-dest-opt-11

Abstract

This document describes an experiment in which VPN service information is encoded in an experimental IPv6 Destination Option. The experimental IPv6 Destination Option is called the VPN Service Option.

One purpose of this experiment is to demonstrate that the VPN Service Option can be deployed in a production network. Another purpose is to demonstrate that the security measures described in this document are sufficient to protect a VPN. Finally, this document encourages replication of the experiment.

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 15 November 2025.

Table of Contents

1. Introduction

Generic Packet Tunneling [RFC2473] allows a router in one network to encapsulate a packet in an IP header and send it to a router in another network. The receiving router removes the outer IP header and forwards the original packet into its own network. This facilitates connectivity between networks that share a private addressing [RFC1918] [RFC4193] plan but are not connected by a direct link.

The IETF refined this concept in a Framework For Virtual Private Networks (VPN) [RFC2764]. It also standardized the following VPN technologies:

* IPSec VPN [RFC3884].

* Layer 3 VPN (L3VPN) [RFC4364].

* Virtual Private LAN Service (VPLS) [RFC4761][RFC4762].

* Layer 2 VPN (L2VPN) [RFC6624].

* Ethernet VPN (EVPN) [RFC7432].

* Pseudowires [RFC8077].

* SRv6 [RFC8986].

* EVPN / NVO3 [RFC9469].

IPSec VPNs cryptographically protect all traffic from customer endpoint to customer endpoint. All of the other VPN technologies mentioned above share the following characteristics:

* An ingress Provider Edge (PE) router encapsulates customer data in a tunnel header. The tunnel header includes service information. Service information identifies a Forwarding Information Base (FIB) entry on an egress PE router.

* The ingress PE router sends the encapsulated packet to the egress PE router.

* The egress PE router receives the encapsulated packet.

* The egress PE router searches its FIB for an entry that matches the incoming service information. If it finds one, it removes the tunnel header and forwards the customer data to a Customer Edge (CE) device, as per the FIB entry. If it does not find a matching FIB entry, it discards the packet.

This document describes an experiment in which VPN service information is encoded in an experimental IPv6 Destination Option [RFC8200]. The experimental IPv6 Destination Option is called the VPN Service Option.

The solution described in this document offers the following benefits:

* It does not require configuration on CE devices.

* It encodes service information in the IPv6 extension header. Therefore, it does not require any non-IPv6 headers (e.g., MPLS) to carry service information.

* It supports many VPNs on a single egress PE router.

* When a single egress PE router supports many VPNs, it does not require an IP address per VPN.

* It does not rely on any particular control plane.

One purpose of this experiment is to demonstrate that the VPN Service Option can be deployed in a production network. Another purpose is to demonstrate that the security measures described in Section 7 of this document are sufficient to protect a VPN. Finally, this document encourages replication of the experiment, so that operational issues can be discovered.

2. Conventions and Definitions

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. The VPN Service Option

The VPN Service Option is an IPv6 Destination Option encoded according to rules defined in [RFC8200].

As described in section 4.2 of [RFC8200] a IPv6 Destination Option contains three fields: Option Type, Opt Data Len, and Option Data. In the VPN Service Option these fields are used as follows:

* Option Type: 8-bit selector. VPN Service Option. This field MUST be set to RFC3692-style Experiment (0x5E) [V6MSG]. See Note below.

* Opt Data Len - 8-bit unsigned integer. Length of the option, in bytes, excluding the Option Type and Option Length fields. This field MUST be set to 4.

* Option Data - 32 bits. VPN Service Information that identifies a FIB entry on the egress PE. The FIB entry determines how the egress PE will forward customer data to a CE device.

A single VPN Service Option MAY appear in a Destination Options header that immediately precedes an upper-layer header. It MUST NOT appear in any other extension header. If a receiver finds the VPN Service Option in any other extension header, it MUST NOT recognize the option. The packet MUST be processed according to the setting of the two highest order bits of the Option Type (see NOTE below).

NOTE: For this experiment, the Option Type is set to ’01011110’, i.e., 0x5E. The highest-order two bits are set to 01 indicating that the required action by a destination node that does not recognize the option is to discard the packet. The third highest-order bit is set to 0 indicating that Option Data cannot be modified along the path between the packet’s source and its destination. The remaining low-order bits are set to ’11110’ to indicate the single IPv6 Destination Option Type code point available in the registry for experimentation.

4. Forwarding Plane Considerations

The ingress PE encapsulates the customer data in a tunnel header. The tunnel header MUST contain an IPv6 header and a Destination Options header that immediately precedes the customer data. It MAY also include any legal combination of IPv6 extension headers.

The IPv6 header contains:

* Version - Defined in [RFC8200]. MUST be equal to 6.

* Traffic Class - Defined in [RFC8200].

* Flow Label - Defined in [RFC8200].

* Payload Length - Defined in [RFC8200].

* Next Header - Defined in [RFC8200].

* Hop Limit - Defined in [RFC8200].

* Source Address - Defined in [RFC8200]. Represents an interface on the ingress PE router. This address SHOULD be chosen according to guidance provided in [RFC6724].

* Destination Address - Defined in [RFC8200]. Represents an interface on the egress PE router. This address SHOULD be chosen according to guidance provided in [RFC6724].

The IPv6 Destination Options Extension Header contains:

* Next Header - Defined in [RFC8200]. MUST identify the protocol of the customer data.

* Hdr Ext Len - Defined in [RFC8200].

* Options - Defined in [RFC8200]. In this experiment, the Options field MUST contain exactly one VPN Service Option as defined in Section 3 of this document. It MAY also contain any legal combination of other Destination Options.

5. Control Plane Considerations

The FIB can be populated:

* By an operator, using a Command Line Interface (CLI).

* By a controller, using the Path Computation Element (PCE) Communication Protocol (PCEP) [RFC5440] or the Network Configuration Protocol (NETCONF) [RFC6241].

* By a routing protocol.

Routing protocol extensions that support the IPv6 VPN Service Destination Option are beyond the scope of this document.

6. IANA Considerations

This document does not make any IANA requests.

7. Security Considerations

A VPN is characterized by the following security policy:

* Nodes outside of a VPN cannot inject traffic into the VPN.

* Nodes inside a VPN cannot send traffic outside of the VPN.

A set of PE routers cooperate to enforce this security policy. If a device outside of that set could impersonate a device inside of the set, it would be possible for that device to subvert security policy. Therefore, impersonation must not be possible. The following paragraphs describe procedures that prevent impersonation.

The IPv6 VPN Service Destination Option can be deployed:

* On the global Internet

* Inside of a limited domain

When IPv6 VPN Service Destination Option is deployed on the global Internet, the tunnel that connects the ingress PE to the egress PE MUST be cryptographically protected by one of the following:

* The IPv6 Authentication Header (AH) [RFC4302]

* The IPv6 Encapsulating Security Payload (ESP) Header [RFC4303].

When IPv6 VPN Service Destination Option is deployed in a limited domain, all nodes at the edge of limited domain MUST maintain Access Control Lists (ACLs). These ACL's MUST discard packets that satisfy the following criteria:

* Contain an IPv6 VPN Service option.

* Contain an IPv6 Destination Address that represents an interface inside of the limited domain.

The mitigation techniques mentioned above operate in fail-open mode. That is, they require explicit configuration in order to ensure that packets using the approach described in this document do not leak out of a domain. See [I-D.wkumari-intarea-safe-limited-domains] for a discussion of fail-open and fail-closed modes.

For further information on the security concerns related to IP tunnels and the recommended mitigation techniques, please see [RFC6169].

8. Deployment Considerations

The VPN Service Option is imposed by an ingress PE and processed by an egress PE. It is not processed by any other nodes along the delivery path between the ingress PE and egress PE.

However, some networks discard packets that include IPv6 Destination Options. This is an impediment to deployment.

Because the VPN Service Option uses an experimental code point, there is a risk of collisions with other experiments. Specifically, the egress PE may process packets from another experiment that uses the same code point.

It is expected that, as with all experiments with IETF protocols, care is taken by the operator to ensure that all nodes participating in an experiment are carefully configured.

Because the VPN Service Destination Option uses an experimental code point, processing of this option MUST be disabled by default. Explicit configuration is required to enable processing of the option.

9. Experimental Results

Parties participating in this experiment should publish experimental results within one year of the publication of this document. Experimental results should address the following:

* Effort required to deploy
  - Was deployment incremental or network-wide?
  - Was there a need to synchronize configurations at each node or could nodes be configured independently?
  - Did the deployment require hardware upgrade?
* Effort required to secure
  - Performance impact
  - Effectiveness of risk mitigation with ACLs
  - Cost of risk mitigation with ACLs
* Mechanism used to populate the FIB
* Scale of deployment
* Interoperability
  - Did you deploy two interoperable implementations?
  - Did you experience interoperability problems?
* Effectiveness and sufficiency of OAM mechanisms
  - Did PING work?
  - Did TRACEROUTE work?
  - Did Wireshark work?
  - Did TCPDUMP work?

10. Acknowledgements

Thanks to Gorry Fairhurst, Antoine Fressancourt, Eliot Lear and Mark Smith for their reviews and contributions to this document.

11. References

11.1. 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/rfc/rfc2119>.

[RFC6169] Krishnan, S., Thaler, D., and J. Hoagland, "Security Concerns with IP Tunneling", RFC 6169, DOI 10.17487/RFC6169, April 2011, <https://www.rfc-editor.org/rfc/rfc6169>.

[RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <https://www.rfc-editor.org/rfc/rfc6724>.

[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/rfc/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/rfc/rfc8200>.

11.2. Informative References

[I-D.wkumari-intarea-safe-limited-domains] Kumari, W., Alston, A., Vyncke, E., Krishnan, S., and D. E. Eastlake, "Safe(r) Limited Domains", Work in Progress, Internet-Draft, draft-wkumari-intarea-safe-limited-domains-04, 3 March 2025, <https://datatracker.ietf.org/doc/html/draft-wkumari-intarea-safe-limited-domains-04>.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <https://www.rfc-editor.org/rfc/rfc1918>.

[RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, DOI 10.17487/RFC2473, December 1998, <https://www.rfc-editor.org/rfc/rfc2473>.

[RFC2764] Gleeson, B., Lin, A., Heinanen, J., Armitage, G., and A. Malis, "A Framework for IP Based Virtual Private Networks", RFC 2764, DOI 10.17487/RFC2764, February 2000, <https://www.rfc-editor.org/rfc/rfc2764>.

[RFC3884] Touch, J., Eggert, L., and Y. Wang, "Use of IPsec Transport Mode for Dynamic Routing", RFC 3884, DOI 10.17487/RFC3884, September 2004, <https://www.rfc-editor.org/rfc/rfc3884>.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <https://www.rfc-editor.org/rfc/rfc4193>.

[RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, December 2005, <https://www.rfc-editor.org/rfc/rfc4302>.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <https://www.rfc-editor.org/rfc/rfc4303>.

[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/rfc/rfc4364>.

[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/rfc/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/rfc/rfc4762>.

[RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <https://www.rfc-editor.org/rfc/rfc5440>.

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <https://www.rfc-editor.org/rfc/rfc6241>.

[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/rfc/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/rfc/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/rfc/rfc8077>.

[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/rfc/rfc8986>.

[RFC9469] Rabadan, J., Ed., Bocci, M., Boutros, S., and A. Sajassi, "Applicability of Ethernet Virtual Private Network (EVPN) to Network Virtualization over Layer 3 (NVO3) Networks", RFC 9469, DOI 10.17487/RFC9469, September 2023, <https://www.rfc-editor.org/rfc/rfc9469>.

[V6MSG] Internet Assigned Numbers Authority (IANA), "Internet Protocol Version 6 (IPv6) Parameters: Destination Options and Hop-by-Hop Options", Web https://www.iana.org/assignments/ipv6-parameters/ipv6-parameters.xhtml#ipv6-parameters-2.

Authors' Addresses

Ron Bonica
Juniper Networks
Herndon, Virginia
United States of America
Email: rbonica@juniper.net

Xing Li
CERNET Center/Tsinghua University
Beijing
China
Email: xing@cernet.edu.cn

Adrian Farrel
Old Dog Consulting
United Kingdom
Email: adrian@olddog.co.uk

Yuji Kamite
NTT Communications Corporation
Minato-ku
Japan
Email: y.kamite@ntt.com

Luay Jalil
Verizon
Richardson, Texas
United States of America
Email: luay.jalil@one.verizon.com