

6man
Internet-Draft
Intended status: Standards Track
Expires: 18 July 2024

R. Bonica
Juniper Networks
Y. Kamite
NTT Communications Corporation
L. Jalil
Verizon
Y. Zhou
ByteDance
G. Chen
Baidu
15 January 2024

The IPv6 Tunnel Payload Forwarding (TPF) Option
draft-bonica-6man-vpn-dest-opt-21

Abstract

This document explains how IPv6 options can be used in IPv6 tunnels. It also defines the IPv6 Tunnel Payload Forwarding (TPF) option.

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 18 July 2024.

Copyright Notice

Copyright (c) 2024 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- 1. Introduction 2
- 2. Requirements Language 4
- 3. The IPv6 Tunnel Payload Forwarding (TPF) Option 4
- 4. TPF Information Determines Next-Protocol Engine Behavior . . 5
- 5. TPF Information Semantics 5
- 6. Virtual Private Networking (VPN) Applications 6
- 7. Security Considerations 6
- 8. IANA Considerations 6
- 9. Acknowledgements 7
- 10. Contributors 7
- 11. References 7
 - 11.1. Normative References 7
 - 11.2. Informative References 8
- Authors' Addresses 8

1. Introduction

This document explains how IPv6 options [RFC8200] can be used in IPv6 tunnels. It also defines the IPv6 Tunnel Payload Forwarding (TPF) option.

An IPv6 tunnel [RFC2473] connects two nodes, called the entry-point and the exit-point. The entry-point receives a packet and encapsulates it in a Tunnel IPv6 Header. Figure 1 depicts the encapsulation.




Figure 1: IPv6 Tunnel Encapsulation

The original packet can be any layer-2 or layer-3 packet (e.g., Ethernet, IPv4, IPv6). The Tunnel Header is an IPv6 header followed by zero or more extension headers. The resulting packet is a Tunnel IPv6 Packet.

The entry-point sends the Tunnel IPv6 Packet to the exit-point which then executes the following procedure:

- * Process the Tunnel IPv6 Header.
- * Remove the Tunnel IPv6 Header, exposing the original packet.
- * Submit the original packet to the next-protocol engine.

The exit-point node processes the Tunnel IPv6 Header in strict left-to-right order. It processes the IPv6 header first and then processes extension headers in the order that they appear in the packet. The IPv6 header, and each extension header, includes a Next Header field. The last Next Header field processed identifies the next-protocol engine.

Entry-point nodes can send optional information to the next-protocol engine on the exit-point node. For example, the entry-point can indicate:

- * The interface through which the next-protocol engine should send the packet.

- * The routing table that the next-protocol engine should use to process the packet.

To send this information, the entry-point node includes an IPv6 Destination Option header in the Tunnel IPv6 Header. The IPv6 Destination Options header includes an IPv6 TPF option and the IPv6 TPF option includes TPF information. The next-protocol engine on the exit-point node uses TPF information when it forwards the original packet.

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. The IPv6 Tunnel Payload Forwarding (TPF) Option

The TPF Option contains the following fields:

- * Option Type: 8-bit selector. TPF option. Value TBD by IANA. (Suggested value: 0x41). See Note below.
- * Opt Data Len - 8-bit unsigned integer. Length of the option, in octets, excluding the Option Type and Option Length fields. This field MUST be set to 4.
- * Option Data - 32-bits. Tunnel Payload Forwarding (TPF) Information.

The TPF option MAY appear in a Destination Options header that precedes an upper-layer header. It MUST NOT appear in a Hop-by-hop Options header or in a Destination Options header that precedes a Routing header.

NOTE : The highest-order two bits of the Option Type (i.e., the "act" bits) are 01. These bits specify the action taken by a destination node that does not recognize the option. The required action is to discard the packet. 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.

4. TPF Information Determines Next-Protocol Engine Behavior

An exit-point node supports one or more next-protocol engines (e.g., Ethernet, IPv4, IPv6). Each next-protocol engine supports a default forwarding procedure and zero or more special forwarding procedures.

When an exit-point node submits a packet to a next-protocol engine without TPF information, the next-protocol engine executes its default forwarding procedure. For example, assume that the exit-point node receives the following Tunnel IPv6 Packet:

- * The Tunnel IPv6 Packet does not contain TPF information.
- * The original packet is IPv4.

In this case, the exit-point node processes and removes the Tunnel IPv6 Header. It then submits the original packet, without any TPF information, to the IPv4 protocol engine.

The IPv4 protocol engine executes its default forwarding procedure. It searches its Forwarding Information Base (FIB) for an entry that matches the original packet's destination address. If the search returns a FIB entry, the protocol engine forwards the packet through an interface that the FIB entry identifies.

When an exit-point node submits a packet to a next-protocol engine with TPF information, the next-protocol engine executes a special forwarding procedure. For example, assume that the exit-point node receives the following Tunnel IPv6 packet:

- * The Tunnel IPv6 Packet contains TPF information that identifies an interface.
- * The original packet is IPv4.

In this case, the exit-point node processes and removes the Tunnel IPv6 Header. It then submits the original packet, along with TPF information, to the IPv4 protocol engine.

The IPv4 protocol engine executes a special forwarding procedure. It forwards the packet through the interface identified by TPF information, without searching the FIB.

5. TPF Information Semantics

TPF information is opaque. While it must be understood by the entry-point node and the exit-point node, it does not need to be understood by any other node.

6. Virtual Private Networking (VPN) Applications

The IPv6 TPF option is useful in deployments where IPv6 tunnels carry:

- * Layer 3 Virtual Private Network (L3VPN) [RFC4364] traffic.
- * Ethernet Virtual Private Network (EVPN) [RFC7432] traffic.

When an IPv6 tunnel carries L3VPN traffic, VPN context information can be encoded in an IPv6 TPF option. Therefore, the MPLS service label that is normally present in an L3VPN packet can be eliminated.

When an IPv6 tunnel carries EVPN traffic, VPN context information can be encoded in an IPv6 TPF option. Therefore, the UDP and VXLAN headers that might otherwise be present can be eliminated.

7. Security Considerations

TPF information MUST NOT be accepted from untrusted sources. The following are acceptable methods of risk mitigation:

- * Authenticate the IPv6 TPF option using the IPv6 Authentication Header (AH) [RFC4302] or the IPv6 Encapsulating Security Payload (ESP) Header [RFC4303].
- * Maintain a secure TPF domain.

All nodes at the edge of a secure TPF domain discard packets that satisfy the following criteria:

- * Contain an IPv6 TPF option.
- * Contain an IPv6 Destination Address that represents an interface inside of the secure TPF domain.

8. IANA Considerations

IANA is requested to allocate a code point 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 "Tunnel Payload Forwarding Option". The "act" bits are 01 and the "chg" bit is 0. The suggested value is 0x41.

9. Acknowledgements

Thanks to Dr. Vanessa Ameen, Brian Carpenter, Adrian Farrel, Ishaan Gandhi, Tom Herbert, John Leddy, Srihari Sangli and Tony Li for their comments.

10. Contributors

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

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/info/rfc2119>>.
- [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/info/rfc2473>>.
- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, December 2005, <<https://www.rfc-editor.org/info/rfc4302>>.

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

11.2. Informative References

- [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>>.
- [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>>.

Authors' Addresses

Ron Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, Virginia 20171
United States of America
Email: rbonica@juniper.net

Yuji Kamite
NTT Communications Corporation
3-4-1 Shibaura, Minato-ku,
108-8118
Japan
Email: y.kamite@ntt.com

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

Yifeng Zhou
ByteDance
Building 1, AVIC Plaza, 43 N 3rd Ring W Rd Haidian District
Beijing
100000
P.R. China
Email: yifeng.zhou@bytedance.com

Gang Chen
Baidu
No.10 Xibeiwang East Road Haidian District
Beijing
100193
P.R. China
Email: phdgang@gmail.com