6man
Internet-Draft
Intended status: Standards Track
Expires: February 11, 2019

R. Bonica
Juniper Networks
J. Leddy
Comcast

August 10, 2018

The IPv6 Probe Option
draft-bonica-6man-unrecognized-opt-03

Abstract

This document defines a new IPv6 option, called the Probe option. The Probe option elicits an ICMPv6 Parameter Problem message from all nodes that process it. When a node sends a packet that contains the Probe option and receives an ICMPv6 Parameter Problem message in response, it has verified the network's ability to convey packets that contain the Probe 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 February 11, 2019.

Table of Contents

1. Introduction

In IPv6 [RFC8200], optional internet-layer information is encoded in extension headers. Two extension headers, the Hop-by-Hop Options header and the Destination Options header, contain a variable number of options. Each option contains the following fields:

- Option Type
- Opt Data Length
- Option Data

The Option Type identifiers are encoded so that their highest-order 2 bits specify the action to be taken if the processing node does not recognize the option. Encodings follow:

- 00 - Skip over the option and continue processing the header.
- 01 - Discard the packet.
- 10 - Discard the packet and send an ICMPv6 [RFC4443] Parameter Problem, Code 2, message to the packet's Source Address, pointing to the unrecognized Option Type.
- 11 - Discard the packet and, only if the packet's Destination Address was not a multicast address, send an ICMPv6 Parameter Problem, Code 2, message to the packet's Source Address, pointing to the unrecognized Option Type.

Several upper-layer protocols [RFC6275] [I-D.leddy-6man-truncate] emit packets that contain IPv6 destination options. These protocols rely the network to convey packets that contain the IPv6 Destination Options header.

A subset of those protocols emit IPv6 destination options with high-order bits equal to "10" and "11". These IPv6 destination options elicit ICMPv6 Parameter Problem messages from destination nodes that do not recognize them. The above-mentioned protocols perform better when the network can convey ICMPv6 Parameter Problem messages from the destination node to the source node.

Operational experience [RFC7872] reveals that a significant number of networks drop all packets that contain the IPv6 Destination Options header. Similarly, a significant number of networks allow packets that contain the IPv6 Destination Options header, but only if Destination Options header does not exceed a specific size. Finally, many networks drop all ICMP Parameter Problem messages.

This document describes procedures by which a source node can discover relevant capabilities of the network that connects it to a destination node. Using these procedures, the source node can determine:

- Whether the network can convey a packet containing a Destination Options header of a specific size from the source node to a destination node.

- Whether the network can convey an ICMPv6 Parameter Problem message from the destination node to the source node.

In order to support the above-mentioned procedures, this document defines a new IPv6 option, called the Probe option. The Probe option elicits an ICMPv6 Parameter Problem message from all nodes that process it. It elicits an IPv6 Parameter Problem message, regardless of whether the processing node recognizes the option. When a source node sends a packet that contains the Probe option and receives an ICMPv6 Parameter Problem message in response, it has verified the above-mentioned network capabilities.

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 Probe Option

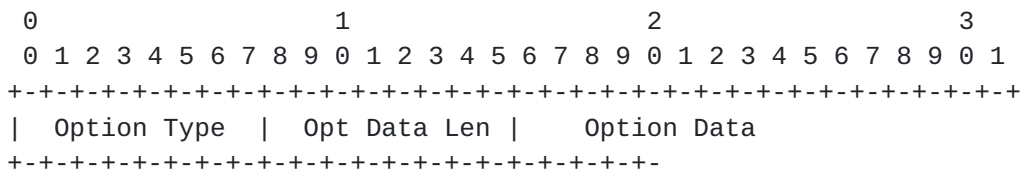

Figure 1

Figure 1 depicts the Probe Option.

Option fields are as follows:

- Option Type - Probe Option. Value TBD by IANA. See Notes below.
- Opt Data Len - Length of Option Data, measured in bytes.
- Option Data - MUST be set to zero on transmission. MUST be ignored on receipt.

The Opt Data Len and Option Data fields can be used to expand the Probe Option and the Destination Options header that contains it to a required length. See Section 4 for details.

A packet MAY contain multiple instances of the Probe option. In IPv6, the maximum size of a Destination Options header is 2048 bytes, while the maximum size of an option instance is only 256 bytes. Therefore, multiple instances of the Probe option are required to expand the Destination Options header beyond 256 bytes.

All nodes process the Probe option as follows, regardless of whether they recognize the option:

- Discard the packet.
- Send an ICMPv6 Parameter Problem, Code 2, message to the packet's Source Address, pointing to the unrecognized Option Type.

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 Probe option. The required action is to discard the packet and send an ICMPv6 Parameter Problem, Code 2, message to the packet's Source Address, pointing to the Probe 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.

4. Discovering Network Capabilities

Assume that a source node needs to determine whether the network can convey a packet from itself to a destination node. The packet contains a Destination Options header whose length is N bytes. As per [RFC8200], the Destination Options header length must be a multiple of 8. Therefore, N must be a multiple of 8.

The source node executes the following procedure:

- Set a short timer (e.g., one or two seconds).
- Send a probe packet.
- Wait for either a) an ICMPv6 Parameter Problem message that matches the probe packet, or b) timer expiration

The probe packet contains an IPv6 Destination Options header and the IPv6 Destination Options header contains one or more instances of Probe option. The number of Probe option instances and the length of Option Data in each instance are chosen so that the Destination Options header length will be equal to N.

In order to influence how the packet is routed to its destination, the probe packet MAY contain upper-layer headers. However, because the packet contains the Probe option, it is always discarded and is never delivered to an upper-layer protocol.

An ICMPv6 Parameter Problem message matches a probe packet if the initial bytes of the probe packet appear in the ICMP Parameter Problem message.

If the source node receives an ICMP Parameter Problem message that matches the probe, both of the following statements are true:

- The network can convey a packet containing a Destination Options header of a specific size from the source node to a destination node.
- The network can convey an ICMPv6 Parameter Problem message from the destination node to the source node.

If the timer expires, at least one of the following statements is true:

- The network cannot convey a packet containing a Destination Options header of a specific size from the source node to a destination node.

- The network cannot convey an ICMPv6 Parameter Problem message from the destination node to the source node.

- Either the probe or the ICMPv6 Parameter Problem message was lost due to a transient issue (e.g., congestion).

As noted above, transient issues can cause false negative results. Therefore, this procedure MAY be repeated after initial failure.

5. Security Considerations

This document introduces no new security vulnerabilities. Any security vulnerabilities exposed by the Probe option are currently exposed by all undefined or unrecognized option types. This is because the Probe option elicits the same behavior as an undefined or unrecognized option

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 "Probe". The "act" bits are 10 and the "chg" bit is 0.

7. Acknowledgements

Thanks to Ross Callon, Fernando Gont and Jinmei Tatuya for their careful review of this document.

8. References

8.1. 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>.

[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.leddy-6man-truncate] Leddy, J. and R. Bonica, "IPv6 Packet Truncation", draft-leddy-6man-truncate-04 (work in progress), June 2018.

[RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <https://www.rfc-editor.org/info/rfc6275>.

[RFC7872] Gont, F., Linkova, J., Chown, T., and W. Liu, "Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World", RFC 7872, DOI 10.17487/RFC7872, June 2016, <https://www.rfc-editor.org/info/rfc7872>.

Authors' Addresses

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

Email: rbonica@juniper.net

John Leddy
Comcast
1717 John F Kennedy Blvd.
Philadelphia, PA 19103
USA

Email: john_leddy@comcast.com