INTERNET-DRAFT
Intended Status: Proposed Standard

N. Elkins
Inside Products
M. Ackermann
BCBS Michigan
K. Haining
US Bank
S. Perdomo
DTCC
W. Jouris
Inside Products
D. Boyes
Sine Nomine

Expires: November 30, 2013

May 30, 2013

IPv6 Performance and Diagnostic Metrics Destination Option
draft-elkins-6man-ipv6-pdm-dest-option-00

Abstract

To diagnose problems for a number of Enterprise Data Center Operators (EDCO), two metrics are critical for timely end-to-end problem resolution, both real-time and after the fact, without impacting an operational production network. They are: packet sequence number and packet timestamp. Packet sequence number is required for diagnostics. Packet timestamp is required to calculate end-to-end response time. Current methods are inadequate for these purposes because they assume unreasonable access to intermediate devices, are cost prohibitive, require infeasible changes to a running production network, and / or do not provide timely data. This document solves these problems with a new destination option, the Performance and Diagnostic Metrics destination option (PDM).

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

Table of Contents

1 Introduction

To diagnose problems for a number of Enterprise Data Center Operators (EDCO), two metrics are critical for timely end-to-end problem resolution, both real-time and after the fact, without impacting an operational production network. They are: packet sequence number and packet timestamp. Packet sequence number is required for diagnostics. Packet timestamp is required to calculate end-to-end response time.

For background, please see Draft-Elkins-Packet-Sequence-Number-Needed-00 [PSNELK], Draft-Elkins-End-To-End-Response-Time-00 [RSPELK], and Draft-Elkins-PDM-Recommended-Usage-00 [USEELK]. These drafts are companion documents to this document. All three documents should be read together.

As discussed in the above Internet Drafts, current methods are inadequate for these purposes because they assume unreasonable access to intermediate devices, are cost prohibitive, require infeasible changes to a running production network, and / or do not provide timely data. This document provides a solution for these problems.

As defined in RFC2460 [RFC2460], destination options are carried by the IPv6 Destination Options extension header. Destination options include optional information that need be examined only by the IPv6 node given as the destination address in the IPv6 header, not by routers in between. This document detail the specifications for a new destination option, the Performance and Diagnostic Metrics destination option (PDM).

1.1 Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2 Performance and Diagnostic Metrics Destination Option

2.1 Destination Options Header

The IPv6 Destination Options Header is used to carry optional information that need be examined only by a packet's destination node(s). The Destination Options Header is identified by a Next Header value of 60 in the immediately preceding header and is defined in RFC2460 [RFC2460].

2.2 Performance and Diagnostic Metrics Destination Option

The IPv6 Performance and Diagnostic Metrics Destination Option (PDM) is an implementation of the Destination Options Header (Next Header value = 60).

It is used to facilitate diagnostics by including a packet sequence number and timestamp.

The PDM destination option is encoded in type-length-value (TLV) format as follows:

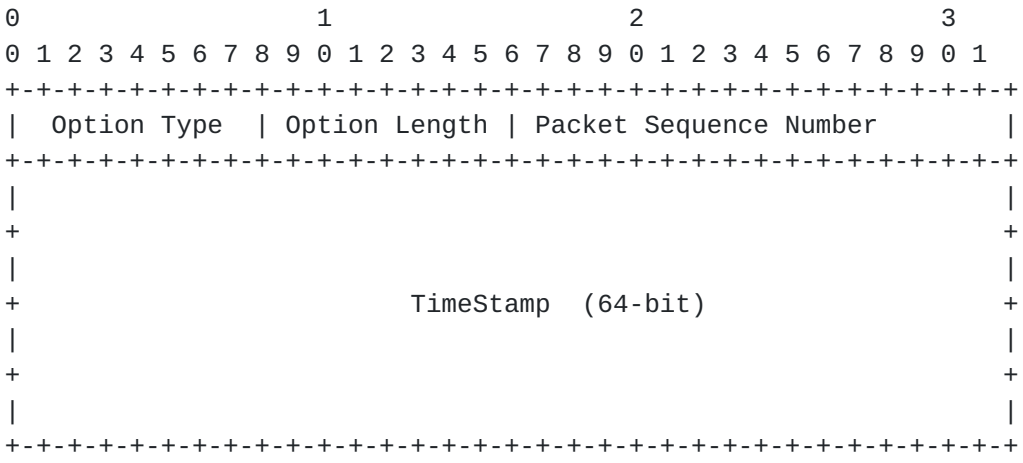

Option Type

TBD = 0xXX (TBD) [To be assigned by IANA] [RFC2780]

Option Length

8-bit unsigned integer. Length of the option, in octets, excluding the Option Type and Option Length fields. This field MUST be set to 80.

Packet Sequence Number

16-bit unsigned integer. This field will wrap. It is intended for human use.

Initialized at 0 and monotonically incremented for protocol packet on the CONNECTION.

Operating systems MUST implement a separate packet sequence number

counter per connection. Operating systems MUST NOT implement a single counter for all connections.

Note: This is consistent with the current implementation of the IPID field in IPv4 for many, but not all, stacks.

TimeStamp

A 64-bit unsigned integer field containing a timestamp. The value indicates the number of seconds since January 1, 1970, 00:00 UTC, by using a fixed point format. In this format, the integer number of seconds is contained in the first 32 bits of the field, and the remaining 32 bits resolve to picoseconds.

This follows timestamp formats used in Network Time Protocol (NTP) [RFC5905] and SEND [RFC3971]. A discussion of why NTP is used in preference to Precision Time Protocol (PTP) is in Draft-Elkins-End-To-End-Response-Time-00.txt.

Implementation note: This format is compatible with the usual representation of time under UNIX, although the number of bits available for the integer and fraction parts in different Unix implementations vary.

Option Type

The two highest-order bits of the Option Type field are encoded to indicate specific processing of the option; for the PDM destination option, these two bits MUST be set to 00. This indicates the following processing requirements:

00 - skip over this option and continue processing the header.

RFC2460 [RFC2460] defines other values for the Option Type field. These MUST NOT be used in the PDM. The other values are as follows:

01 - discard the packet.

10 - discard the packet and, regardless of whether or not the packet's Destination Address was a multicast address, send an ICMP 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 ICMP Parameter Problem, Code 2, message to the packet's Source Address, pointing to the unrecognized Option Type.

In keeping with RFC2460 [RFC2460], the third-highest-order bit of the Option Type specifies whether or not the Option Data of that option can change en-route to the packet's final destination.

In the PDM, the value of the third-highest-order bit MUST be 0. The possible values are as follows:

0 - Option Data does not change en-route

1 - Option Data may change en-route

The three high-order bits described above are to be treated as part of the Option Type, not independent of the Option Type. That is, a particular option is identified by a full 8-bit Option Type, not just the low-order 5 bits of an Option Type.

Header Placement

The PDM destination option MUST be placed as follows:

- Before the upper-layer header. That is, this is the last extension header.

This follows the order defined in RFC2460 [RFC2460]

IPv6 header

Hop-by-Hop Options header

Destination Options header

Routing header

Fragment header

Authentication header

Encapsulating Security Payload header

Destination Options header

upper-layer header

For each IPv6 packet header, the PDM MUST NOT appear more than once. However, an encapsulated packet MAY contain a separate PDM associated with each encapsulated IPv6 header.

The inclusion of a PDM in a packet affects the receiving node's processing of only this single packet. No state is created or modified in the receiving node as a result of receiving a PDM in a packet.

2.3 Implementation Considerations

The PDM destination options extension header SHOULD be turned on by each stack on a host node.

If implemented, each operating system MUST have a default configuration parameter, e.g. diag_header_sys_default_value=yes/no. The operating system MAY also have a dynamic configuration option to change the configuration setting as needed.

If the PDM destination options extension header is used, then it MAY be turned on for all packets flowing through the host, applied to an upper-layer protocol (TCP, UDP, SCTP, etc), a local port, or IP address only. These are at the discretion of the implementation.

The PDM MUST NOT be changed dynamically via packet flow as this may create potential security violation or DoS attack by numerous packets turning the header on and off.

As with all other destination options extension headers, the PDM is for destination nodes only. As specified above, intermediate devices MUST neither set nor modify this field.

3 Backward Compatibility

The scheme proposed in this document is backward compatible with all the currently defined IPv6 extension headers. According to RFC2460 [RFC2460], if the destination node does not recognize this option, it should skip over this option and continue processing the header.

4 Security Considerations

The PDM MUST NOT be changed dynamically via packet flow as this creates a possibility for potential security violations or DoS attacks by numerous packets turning the header on and off.

5 IANA Considerations

An option type must be assigned by IANA for the Performance and Diagnostic Metrics destination option.

6 References

6.1 Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[PSNELK] Elkins, N., "Draft-Elkins-Packet-Sequence-Number-Needed-00", Internet Draft, May 2013.

[RSPELK] Elkins, N., "Draft-Elkins-End-To-End-Response-Time-00", Internet Draft, May 2013

[USEELK] Elkins, N., "Draft-Elkins-PDM-Recommended-Usage-00", Internet Draft, May 2013

6.2 Informative References

[RFC2780] Bradner, S. and V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", BCP 37, RFC 2780, March 2000.

[RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

[RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, June 2010.

7 Acknowledgments

The authors would like to thank Rick Troth and Fred Baker for their comments.

Authors' Addresses

Nalini Elkins
Inside Products, Inc.
36A Upper Circle
Carmel Valley, CA 93924
United States
Phone: +1 831 659 8360
Email: nalini.elkins@insidethestack.com
http://www.insidethestack.com

Michael S. Ackermann
Blue Cross Blue Shield of Michigan
P.O. Box 2888
Detroit, Michigan 48231
United States
Phone: +1 310 460 4080
Email: mackermann@bcbsmi.com
http://www.bcbsmi.com

Keven Haining
US Bank
16900 W Capitol Drive
Brookfield, WI 53005
United States
Phone: +1 262 790 3551
Email: keven.haining@usbank.com
http://www.usbank.com

Sigfrido Perdomo
Depository Trust and Clearing Corporation
55 Water Street
New York, NY 10055
United States
Phone: +1 917 842 7375
Email: s.perdomo@dtcc.com
http://www.dtcc.com

William Jouris
Inside Products, Inc.
36A Upper Circle
Carmel Valley, CA 93924
United States
Phone: +1 925 855 9512
Email: bill.jouris@insidethestack.com
http://www.insidethestack.com

David Boyes
Sine Nomine Associates
43596 Blacksmith Square
Ashburn, VA 20147
United States
Phone: +1 703 723 6673
dboyes@sinenomine.net
http://www.sinenomine.net