

INTERNET-DRAFT

N. Elkins
B. Jouris
Inside Products
M. Ackermann
BCBS Michigan
July 29, 2014

Intended Status: Proposed Standard
Expires: January 2015

**IPv6 Performance and Diagnostic Metrics Destination Option
draft-elkins-6man-ipv6-pdm-dest-option-06**

Abstract

To diagnose performance and connectivity problems, metrics on passive traffic are critical for timely end-to-end problem resolution. Such diagnostics may be real-time or after the fact, but must not impact an operational production network. The base metrics are: packet sequence number and packet timestamp. Metrics derived from these will be described separately. This document solves these problems with a new destination option, the Performance and Diagnostic Metrics destination option (PDM).

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/1id-abstracts.html>

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

Copyright and License Notice

Copyright (c) 2014 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 (<http://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	3
1.1	Terminology	3
2	Performance and Diagnostic Metrics Destination Option	3
2.1	Destination Options Header	3
2.2	PDM Types	3
2.3	Performance and Diagnostic Metrics Destination Option	3
2.4	Header Placement	7
2.5	Implementation Considerations	8
2.5.1	Dynamic Configuration Options	8
2.5.2	Data Length Filtering	8
2.5.3	5-tuple Aging	8
3	Backward Compatibility	9
4	Security Considerations	9
5	IANA Considerations	9
6	References	9
6.2	Informative References	10
7	Acknowledgments	10
	Authors' Addresses	10

1 Introduction

To diagnose performance and connectivity problems, metrics on passive traffic are critical for timely end-to-end problem resolution. Such diagnostics may be real-time or after the fact, but must not impact an operational production network. The base metrics are: packet sequence number and packet timestamp.

For background, please see [draft-zheng-framework-passive-01](#) [Passive]. This draft provides a foundation to this document.

As discussed in the above Internet Draft, passive measurement is preferred for measuring end-user Quality of Experience (QoE) as well as peak capacity demands. This document provides a solution for obtaining such metrics.

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 or other "middle boxes". This document specifies 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 PDM Types

The IPv6 Performance and Diagnostic Metrics Destination Option (PDM) is an implementation of the Destination Options Header (Next Header value = 60). The PDM does not require time synchronization.

2.3 Performance and Diagnostic Metrics Destination Option

Elkins

Expires January 30, 2015

[Page 3]

The Performance and Diagnostic Metrics Destination Option (PDM) contains the following fields:

PSNTP : Packet Sequence Number This Packet
 PSNLR : Packet Sequence Number Last Received
 DELTALR : Delta Last Received
 PSNLS : Packet Sequence Number Last Sent
 DELTALS : Delta Last Sent

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

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Option Type | Option Length | PSN This Packet |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| PSN Last Received | PSN Last Sent |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Delta Last Received | Delta Last Sent |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| TType |
+--+--+--+--+

```

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 22.

Packet Sequence Number This Packet (PSNTP)

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

Initialized at a random number and monotonically incremented for packet on the 5-tuple. The 5-tuple consists of the source and destination IP addresses, the source and destination ports, and the

Elkins

Expires January 30, 2015

[Page 4]

upper layer protocol (ex. TCP, ICMP, etc).

Operating systems MUST implement a separate packet sequence number counter per 5-tuple. 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.

Packet Sequence Number Last Received (PSNLR)

16-bit unsigned integer. This is the PSN of the packet last received on the 5-tuple.

Packet Sequence Number Last Sent (PSNLS)

16-bit unsigned integer. This is the PSN of the packet last sent on the 5-tuple.

Delta TimeStamp Type (TIMETYPE)

4-bit unsigned integer. This is the type of time contained in the delta fields below.

0 - unknown 1 - time is in units of nanoseconds 2 - time is in units of microseconds 3 - time is in units of milliseconds 4 - time is in units of seconds 5 - time is in units of minutes 6 - time is in units of hours 7 - time is in units of days

The values 5 - 7 are relevant for Delay Tolerant Networks (DTN) which may operate with long delays between packets.

Delta Last Received (DELTALR)

A 16-bit unsigned integer field. This is server delay.

$DELTALR = \text{Send time packet 2} - \text{Receive time packet 1}$

The value is according to the scale in TIMETYPE.

Delta Last Sent (DELTALS)

A 16-bit unsigned integer field. This is round trip or end-to-end

time.

Delta Last Sent = Receive time packet 2 - Send time packet 1

The value is in according to the scale in TIMETYPE.

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.

2.4 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.

2.5 Implementation Considerations

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

2.5.1 Dynamic Configuration Options

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.

2.5.2 Data Length Filtering

Different results for derived metrics, such as, server delay, will be obtained if calculations are done including or excluding packets which have a data length of 0 or 1. Some protocols, for example, TCP, provide acknowledgements which have a length of 0. Keep-alive packets have a data length of 0 or 1.

Operating systems may provide the user a choice of whether to include or exclude packets with a 0 or 1 byte data length.

2.5.3 5-tuple Aging

Within the operating system, metrics must be kept on a 5-tuple basis. As discussed before, these are:

- PSNTP : Packet Sequence Number This Packet
- TSTP : Timestamp This Packet
- PSNLR : Packet Sequence Number Last Received
- TSLR : Timestamp Last Received
- PROTC : Protocol for Upper Layer (ex. TCP, UDP, ICMP, etc)

The question comes of when to stop keeping data or restarting the

Elkins

Expires January 30, 2015

[Page 8]

numbering for a 5-tuple. For example, in the case of TCP, at some point, the connection will terminate. Keeping data in control blocks forever, will have unfortunate consequences for the operating system.

So, the recommendation is to use a known aging parameter such as Max Segment Lifetime (MSL) as defined in Transmission Control Protocol [RFC0793]. The choice of aging parameter is left up to the implementation.

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

- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.
- [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.
- [RFC2780] Bradner, S. and V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", [BCP 37](#), [RFC 2780](#), March 2000.

6.2 Informative References

[Passive] Zheng, V., "[draft-zheng-ippm-framework-passive-01](#)", Internet Draft, July 2014. [Work in Progress]

7 Acknowledgments

The authors would like to thank Keven Haining, Sigfrido Perdomo 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>

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>

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>

