

INTERNET-DRAFT
Intended Status: Informational

N. Elkins
Inside Products
M. Ackermann
BCBS Michigan
K. Haining
US Bank
S. Perdomo
DTCC
W. Jouris
Inside Products
D. Boyes
Sine Nomine
May 30, 2013

Expires: November 30, 2013

End-to-end Response Time Needed for IPv6 Diagnostics
draft-elkins-v6ops-ipv6-end-to-end-rt-needed-00

Abstract

For a number of Enterprise Data Center Operators (EDCO) both real-time and after the fact problem resolution is critical. Two metrics are critical for timely end-to-end problem resolution, 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, or do not provide timely data. This document provides the background and rationale for the packet timestamp which is a part of the IPv6 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/lid-abstracts.html>

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

INTERNET DRAFT -elkins-v6ops-ipv6-end-to-end-rt-needed-00

May 2013

Copyright and License Notice

Copyright (c) 2013 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	Background	3
1.1	Why End-to-end Response Time is Needed	3
1.2	Trending of Response Time Data	4
1.3	What to measure?	4
1.4	TCP Timestamp not enough	5
1.5	Inadequacy of Current Instrumentation Technology	5
1.5.1	Synthetic transactions	5
1.5.2	PING	5
1.5.3	Other Estimates of Network Time	6
1.5.4	Server / Client Agents	6
2	Solution Parameters	6
2.1	Rationale for proposed solution	7
2.2	Merits of timestamp in PDM	7
2.3	What kind of timestamp?	8
3	Backward Compatibility	8
4	Security Considerations	8
5	IANA Considerations	8
6	References	8
6.1	Informative References	8
7	Acknowledgments	9

[1](#) Background

For a number of Enterprise Data Center Operators (EDCO) both real-time and after the fact problem resolution is critical. Two metrics are critical for timely end-to-end problem resolution, 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.

This document provides the background and rationale for the packet timestamp which is a part of the IPv6 Performance and Diagnostic Metrics destination option (PDM).

For background, please see Draft-Elkins-IPv6-PDM-Dest-Option-00 [[PDMELK](#)], Draft-Elkins-Packet-Sequence-Number-Needed-00 [[PSNELK](#)], and Draft-Elkins-PDM-Recommended-Usage-00 [[USEELK](#)]. These drafts are companion documents to this document. All four 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, or do not provide timely data. The IPv6 Performance and Diagnostic Metrics destination option (PDM) provides a solution to these problems. This document will detail the background and need for the packet sequence number.

[1.1](#) Why End-to-end Response Time is Needed

The timestamp traveling along with the packet will be used to calculate end-to-end response time, without requiring agents in

devices along the path. In many EDCO Networks, end-to-end response times are a critical component of Service Levels Agreements (SLAs). End-to-end response is what the user of a network system actually experiences. When the end user is an individual, he is generally indifferent to what is happening along the network; what he really cares about is how long it takes to get a response back. But this is not just a matter of individuals' personal convenience. In many cases, rapid response is critical to the business being conducted.

When the end user is a host (e.g. when doing a file transfer), what matters is the speed with which requested data can be transferred -- specifically, whether the requested data can be transferred in time to accomplish the desired actions. This can be important when the relevant external conditions are subject to rapid change.

Response time and consistency are not just "nice to have". On many

Elkins

Expires November 15, 2013

[Page 3]

INTERNET DRAFT -elkins-v6ops-ipv6-end-to-end-rt-needed-00

May 2013

networks, the impact can be financial hardship or endanger human life. In some cities, the emergency police contact system operates over IP, law enforcement uses TCP/IP networks, our stock exchanges are settled using IP networks. The critical nature of such activities to our daily lives and financial well-being demand a solution. [Section 1.5](#) will detail the current state of end-to-end response time monitoring today.

[1.2](#) Trending of Response Time Data

In addition to the need for tracking current service, end-to-end response time is valuable for capacity planning. By tracking response times, and identifying trends, it becomes possible to determine when network capacity is being approached. This allows additional capacity to be obtained before service levels fall below requirements. Without that kind of tracking, the only option is to wait until there is a problem, and then scramble to get additional capacity on an emergency (and probably high cost) basis.

The document Draft-Elkins-PDM-Recommended-Usage-00 [[USEELK](#)] will detail our recommended use for the PDM for capacity planning purposes.

[1.3](#) What to measure?

End to end response time can be broken down into 3 parts:

- Network time
- Application (or server) time
- Client time

Of course, network time may include multiple hops, application and server time may be exacerbated by stack time. By and large, the three timings are 'good enough' measurements to allow rapid triage into the failing component.

Ways are available (provided by operating systems) to measure Application and Client times. Network time can also be measured in isolation via some of the measurement techniques described in [section 1.5](#). The most difficult portion is to integrate network time with the server or application times. Products exist to do this but are available at an exorbitant cost, require agents, and will likely become more prohibitive as the speed of networks grow and as the world becomes more connected via mobile devices. This is discussed in detail in [section 1.5](#).

Measuring network time requires precise timestamps. Furthermore, those timestamps need to occur at the end-points of the transactions

being measured. And they need to be available, regardless of the protocol being used by the transaction. Which is to say, the timestamp has to be available in one of the extensions to the IP header - this is provided by the PDM.

[1.4](#) TCP Timestamp not enough

Some suggest that the TCP Timestamp option might be sufficient to calculate end-to-end response time.

The TCP Timestamp Option is defined in [RFC1323](#) [[RFC1323](#)]. The reason for the TCP Timestamp option is to be able to discard packets when the TCP sequence number wraps. (PAWS)

The problems with the TCP Timestamp option are:

1. Not everyone turns this on.

2. It is only available for TCP applications
3. No time synchronization between sender and receiver.
4. No indication of date in long-running connections. (That is connections which last longer than one day)
5. The granularity of the timestamp is at best at millisecond level. In the future, as speeds of devices and networks grow, this level of granularity will be inadequate. Even today, on many networks, the timings are at microsecond level not millisecond.

[1.5](#) Inadequacy of Current Instrumentation Technology

The current technology includes:

1. Synthetic transactions
2. Pings
3. Other Estimates of network time
4. Server / Client Agents

[1.5.1](#) Synthetic transactions

[1.5.2](#) PING An ICMP ping measures network time. First, you can PING the remote device. Then you assume that the time it takes to get a response to a PING is the same as the time that a transaction (regardless of packet size) would take to traverse the network.

However, QoS rules, firewalls, etc. may mean that PING, (and other synthetic transactions) may not be subject to the same conditions.

[1.5.3](#) Other Estimates of Network Time

If a packet trace is done, it is possible to look at the time between when a response was seen to be sent at the packet capture device and when the ACK for the response comes back.

If you assume that the ACK took the same amount of time as the original query, you have the network time. Unfortunately, the time

for the ACK may not be the same as the time for a much larger query transaction to traverse the network.

The biggest problem with this method is that of TCP delayed acknowledgements. If the client is doing delayed ACKs, then the ACK will be held until the next request is ready to go out. In this case, the time to receive the ACK has no correlation with network time.

[1.5.4](#) Server / Client Agents

There are also products which claim that they can determine end-to-end response times, integrating server and network times - and indeed they can do so. But they require agents which must be placed at each point which is to be monitored. That is, it is necessary to add those agents EVERYWHERE around the network, at a very high cost. These kind of products can be purchased by only the richest 1% of the corporations. As the speed of networks grow, and as the world becomes more connected via mobile devices, such products will only become more expensive. If, indeed, their technology can keep up.

TCP/IP networks today are used throughout the world. The need for adequate performance will become more and more critical. A method that is scalable and affordable is needed to ensure this growth.

[2](#) Solution Parameters

What is needed is:

- 1) A method to identify and/or track the behavior of a connection without assuming access to the transport devices.
- 2) A method to observe a connection in flight without introducing agents.
- 3) a method to observe arbitrary flows at multiple points within a network and correlate the results of those observations in a

consistent manner.

- 4) A method to signal and correlate transport issues to application end-to-end behavior.

- 5) A method which does not require changes to a production network in real time.
- 6) Adequate granularity in the measurement technique to provide the needed metrics.
- 7) A method that is scalable to very large networks.
- 8) A method that is affordable to all.

2.1 Rationale for proposed solution

The current IPv6 specification does not provide a timestamp number nor similar field in the IPv6 main header or in any extension header. So, we propose the IPv6 Performance and Diagnostic Metrics destination option (PDM) [[PDMELK](#)].

2.2 Merits of timestamp in PDM

Advantages include:

1. Less overhead than other alternatives.
2. Real measure of actual transactions.
3. Less cost to provide solutions
4. More accurate and complete information.
5. Independence from transport layer protocols.
6. Ability to span organizational boundaries with consistent instrumentation

In other words, this is a solution to a long-standing problem. The PDM will provide a metric which will allow those responsible for network support to determine what is happening in their network without expensive equipment (agents) at each device.

The PDM does not solve every response time issue for every situation. Network connections with multiple hops will still need more granular metrics, as will the differentiation between multiple components at each host. That is, TCP/IP stack time vs. applications time will

still need to be broken out by client software. What the PDM does provide is triage. That is, to determine quickly if the problem is in the network or in the server or application.

[2.3](#) What kind of timestamp?

Questions arise about exactly the kind of timestamp to use. Both the Network Time Protocol (NTP) [[RFC5905](#)] and Precision Time Protocol (PTP) [[IEEE1588](#)] are used to provide timing on TCP/IP networks.

NTP has evolved within the IETF structure while PTP has evolved within the Institute of Electrical and Electronics Engineers (IEEE) community. By and large, operating systems such as Windows, Linux, and IBM mainframe computers use NTP. These are the source and destination systems for packets. Intermediate nodes such as routers and switches may prefer PTP.

Since we are describing a new extension header for destination systems, the timestamp to be used will be in accordance with NTP. Many, if not all, EDCO networks currently use NTP. In the document Draft-Elkins-PDM-Recommended-Usage-00 [[USEELK](#)], we will discuss our recommendations for NTP and the PDM.

[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

There are no security considerations.

[5](#) IANA Considerations

There are no IANA considerations.

[6](#) References

[6.1](#) Informative References

[RFC1323] Jacobson, V., Braden, R., and D. Borman, "TCP Extensions for High Performance", [RFC 1323](#), May 1992.

INTERNET DRAFT -elkins-v6ops-ipv6-end-to-end-rt-needed-00

May 2013

- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", [RFC 5905](#), June 2010.
- [IEEE1588] IEEE 1588-2002 standard, "Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems"
- [PDMELK] Elkins, N., "Draft-Elkins-IPv6-PDM-Dest-Option-00", Internet Draft, May 2013.
- [PSNELK] Elkins, N., "Draft-Elkins-Packet-Sequence-Number-Needed-00", Internet Draft, May 2013.
- [USEELK] Elkins, N., "Draft-Elkins-PDM-Recommended-Usage-00", Internet Draft, May 2013

[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.insidestack.com>

Michael S. Ackermann
Blue Cross Blue Shield of Michigan
P.O. Box 2888
Detroit, Michigan 48231
United States
Phone: +1 310 460 4080

Elkins

Expires November 15, 2013

[Page 9]

INTERNET DRAFT -elkins-v6ops-ipv6-end-to-end-rt-needed-00

May 2013

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@insidestack.com
<http://www.insidestack.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>

Elkins

Expires November 15, 2013

[Page 10]