

Network Working Group
Internet-Draft
Updates: [2330](#) (if approved)
Intended status: Informational
Expires: April 16, 2016

J. Fabini
TU Wien
A. Morton
AT&T Labs
October 14, 2015

Updates for IPPM's Framework: Timestamping and Use Cases
draft-fabini-ippm-2330-time-00

Abstract

Quality and accuracy of measurements depend on the selection of appropriate locations and timers for timestamp acquisition. This memo updates the IP Performance Metrics (IPPM) Framework [RFC 2330](#) with new considerations on timers, timestamps and time-related definitions with particular focus on wireless networks and virtualized hosts.

Requirements Language

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](#)].

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 <http://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 April 16, 2016.

Copyright Notice

Copyright (c) 2015 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	2
2.	Scope	3
3.	Definition and Use of Wire Time	4
3.1.	Virtualization	5
4.	Definition and Use of Host Time	5
4.1.	Virtualization	6
5.	Encryption	6
6.	Security Considerations	6
7.	IANA Considerations	7
8.	Acknowledgements	7
9.	References	7
9.1.	Normative References	7
9.2.	Informative References	9
	Authors' Addresses	9

[1.](#) Introduction

The IETF IP Performance Metrics (IPPM) working group first created a framework for metric development in [[RFC2330](#)]. This framework has stood the test of time and enabled development of many fundamental metrics. It has been updated in the area of metric composition [[RFC5835](#)], and in several areas related to active stream measurement of modern networks with reactive properties [[RFC7312](#)].

Time is fundamental to measurements. The IPPM framework [[RFC2330](#)] includes a comprehensive terminology, definition and discussion of time- and clock-related concepts. But network technology has evolved since 1998. First, host computing performance and software complexity have increased while network delays have decreased. These advances enable development of new concepts like host virtualization in cloud computing, and network function virtualization. One side-effect of these technologies is that packets spend more and more time in software, their propagation through the stack being biased by a series of uncertainty factors like buffers, queues and operating system scheduling. Measurements can acquire timestamps in various locations within the host's software stack, causing potentially

significant variances in timestamp values [[MCDC](#)]. This is a consequence of just one alternative that the IPPM framework [[RFC2330](#)] defines for these timestamps: host time. This memo seeks to update this aspect of the IPPM framework by discussing alternatives and proposing generic solutions that can be referenced by metric definitions.

Second, the IPPM framework was approved one year before the release of the first IEEE 802.11 series standard. This was a time when the data communication landscape was dominated by wired communication technologies and wireless data communications were in their infancy. This fact originated IPPM definitions like host time or wire time. Today, after more than one decade and half, wireless IEEE 802.11 networks and cellular packet-switched data technologies like High-Speed Packet Access (HSPA) and Long Term Evolution (LTE) are responsible for an essential part of the worldwide Internet access traffic. This raises the question, how the IPPM framework concept of wire time can be mapped to wireless networks.

This memo revisits and updates time-related IPPM framework [[RFC2330](#)] definitions like host time and wire time in the light of technological advances.

2. Scope

The purpose of this memo is to update the IPPM metrics framework [[RFC2330](#)] with respect to timestamps and virtualization. Concepts like "host time" or "wire time" must be revisited to be applicable for wireless networks and virtualized environments.

The scope is to update key sections of [[RFC2330](#)] ...

Potential Topics (Draft outline for discussion):

- Fundamentals

1. Host Time (further develop the definition as mentioned but not defined in [RFC2330](#). Explain variants, propose examples and guidance). Dedicated subsection and challenge: virtualization.

2. Wire Time (rename "Wire Time" to "Medium Time" or "Media Time" to account for wireless networks; revisit the definition of "Media Arrival Time" and "Media Exit Time" in the light of encryption, complex network coding, interleaving etc. Define how to measure media time in wireless networks) Dedicated subsection and challenge: virtualization

3. Define new terms wrt timestamps that help to differentiate between software delays within a host. Ideally it should give guidance and define terms that allow to differentiate between delays (uncertainties) that a) are caused by host load, timer resolution, system scheduling, etc. and b) the ones that originate when packets wait in drivers because of network access policies or network congestion. A "network (or IP) timestamp" can be useful in (virtualized) environments it might make sense to define the term "network timestamp" as "tcpdump timestamp of the host instance that is located closest to the hardware". Hypervisor tcpdump timestamps should be preferred over VM timestamps.
4. (optional) Protocol design considerations: (Security vs. flexibility): adding timestamps into headers/payload protocol fields at driver level (tcpdump equivalent) is difficult/impossible if link is protected using SSL/TLS/IPsec. Meaningful: insert timestamp in application space.
5. (probably not in IPPM) NTP is understood to operate sub-optimally for time-transfer to virtual machines (VM) as guest-clients in some hosts. What mechanisms can be employed to improve time accuracy and stability when VMs intend to perform time measurement?

3. Definition and Use of Wire Time

[RFC2330] defines the following terms related to wire time (defined in terms of an Internet host H observing an Internet link L at a particular location):

- o For a given packet P, the 'wire arrival time' of P at H on L is the first time T at which any bit of P has appeared at H's observational position on L
- o For a given packet P, the 'wire exit time' of P at H on L is the first time T at which all the bits of P have appeared at H's observational position on L.


However, wireless LAN and cellular networks are essential components of today's Internet. This memo proposes to replace the term "wire time" by the equivalent term "media time" and provides additional guidance for wireless network measurements.

This definition of media time raises one additional challenge: how can "media arrival time" and "media exit time" be defined for wireless networks. Aligning with existing physical-layer standards of other standardization organizations like the 3GPP, this memo recommends the antenna connector of the Host(?) to be used for measuring media time in wireless networks.

3.1. Virtualization

Virtualization adds additional level of uncertainty for timestamps: VM application, VM tcpdump, Hypervisor tcpdump. What about wire time for virtual network interfaces? Should it be bound to the related physical interface or to the hypervisor/virtual interface?

There was some flexibility in the original mention of Host Time. Virtualization adds more layers where host timestamps could be applied or derived, and the specification requires its own framework:



4. Definition and Use of Host Time

[Section 3 of \[RFC2330\]](#) includes an occurrence of the term "host time" but does not define it. Some metric definitions like one-way delay in [section 3.7.2 of \[RFC2679\]](#) or round-trip delay in [section 2.7 of \[RFC2681\]](#) discuss the consequences of differences between wire time and host time. Common to these metric definitions is the request to report errors and uncertainties resulting from the difference between host time and wire time.

However, the evolution of network technology has decreased the network delay substantially. Today, the delay of wireless and wired LAN links are in the same order of magnitude as the uncertainty of host time, the difference between timestamps at application layer and the ones in kernel space being substantial as pointed out in [\[MCD\]](#)). Instead of accepting this uncertainty as an inherent part of measurements, this memo recommends to accurately define the location within the stack where timestamps are assigned to a packet.

Virtualization extends the concept of host time beyond earlier limits. For instance packet captures (tcpdump) in virtualized environments are possible at VM-level and at hypervisor level. Therefore, the definition of "Host Time" spans a broad range of timestamps, encompassing anything from timestamps acquired by an application running on top of a VM, down to tcpdump timestamps acquired within the VM or hypervisor network stack.

<http://www.cc.gatech.edu/~lingliu/papers/2012/YiRen-CollaborateCom2012.pdf>

The security considerations that apply to any active measurement of live paths are relevant here as well. See [RFC4656] and [RFC5357].

When considering privacy of those involved in measurement or those whose traffic is measured, the sensitive information available to potential observers is greatly reduced when using active techniques which are within this scope of work. Passive observations of user traffic for measurement purposes raise many privacy issues. We refer the reader to the privacy considerations described in the Large Scale Measurement of Broadband Performance (LMAP) Framework [[RFC7594](#)], which covers active and passive techniques.

7. IANA Considerations

This memo makes no requests of IANA.

8. Acknowledgements

The authors thank ...

9. References

9.1. Normative References

- [RFC0791] Postel, J., "Internet Protocol", STD 5, [RFC 791](#), DOI 10.17487/RFC0791, September 1981, <<http://www.rfc-editor.org/info/rfc791>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", [RFC 2330](#), DOI 10.17487/RFC2330, May 1998, <<http://www.rfc-editor.org/info/rfc2330>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), DOI 10.17487/RFC2460, December 1998, <<http://www.rfc-editor.org/info/rfc2460>>.
- [RFC2675] Borman, D., Deering, S., and R. Hinden, "IPv6 Jumbograms", [RFC 2675](#), DOI 10.17487/RFC2675, August 1999, <<http://www.rfc-editor.org/info/rfc2675>>.
- [RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", [RFC 2679](#), DOI 10.17487/RFC2679, September 1999, <<http://www.rfc-editor.org/info/rfc2679>>.

- [RFC2681] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM", [RFC 2681](#), DOI 10.17487/RFC2681, September 1999, <<http://www.rfc-editor.org/info/rfc2681>>.
- [RFC2780] Bradner, S. and V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", [BCP 37](#), [RFC 2780](#), DOI 10.17487/RFC2780, March 2000, <<http://www.rfc-editor.org/info/rfc2780>>.
- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", [RFC 3168](#), DOI 10.17487/RFC3168, September 2001, <<http://www.rfc-editor.org/info/rfc3168>>.
- [RFC4494] Song, JH., Poovendran, R., and J. Lee, "The AES-CMAC-96 Algorithm and Its Use with IPsec", [RFC 4494](#), DOI 10.17487/RFC4494, June 2006, <<http://www.rfc-editor.org/info/rfc4494>>.
- [RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", [RFC 4656](#), DOI 10.17487/RFC4656, September 2006, <<http://www.rfc-editor.org/info/rfc4656>>.
- [RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", [RFC 5357](#), DOI 10.17487/RFC5357, October 2008, <<http://www.rfc-editor.org/info/rfc5357>>.
- [RFC5644] Stephan, E., Liang, L., and A. Morton, "IP Performance Metrics (IPPM): Spatial and Multicast", [RFC 5644](#), DOI 10.17487/RFC5644, October 2009, <<http://www.rfc-editor.org/info/rfc5644>>.
- [RFC5835] Morton, A., Ed. and S. Van den Berghe, Ed., "Framework for Metric Composition", [RFC 5835](#), DOI 10.17487/RFC5835, April 2010, <<http://www.rfc-editor.org/info/rfc5835>>.
- [RFC6282] Hui, J., Ed. and P. Thubert, "Compression Format for IPv6 Datagrams over IEEE 802.15.4-Based Networks", [RFC 6282](#), DOI 10.17487/RFC6282, September 2011, <<http://www.rfc-editor.org/info/rfc6282>>.
- [RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", [RFC 6437](#), DOI 10.17487/RFC6437, November 2011, <<http://www.rfc-editor.org/info/rfc6437>>.

- [RFC6564] Krishnan, S., Woodyatt, J., Kline, E., Hoagland, J., and M. Bhatia, "A Uniform Format for IPv6 Extension Headers", [RFC 6564](#), DOI 10.17487/RFC6564, April 2012, <<http://www.rfc-editor.org/info/rfc6564>>.
- [RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", [RFC 7045](#), DOI 10.17487/RFC7045, December 2013, <<http://www.rfc-editor.org/info/rfc7045>>.
- [RFC7312] Fabini, J. and A. Morton, "Advanced Stream and Sampling Framework for IP Performance Metrics (IPPM)", [RFC 7312](#), DOI 10.17487/RFC7312, August 2014, <<http://www.rfc-editor.org/info/rfc7312>>.

9.2. Informative References

- [MCDC] Fabini, J. and T. Zseby, "M2M communication delay challenges: Application and measurement perspectives", IEEE International Instrumentation and Measurement Technology Conference (I2MTC 2015) doi: 10.1109/I2MTC.2015.7151564, May 2015.
- [RFC7594] Eardley, P., Morton, A., Bagnulo, M., Burbridge, T., Aitken, P., and A. Akhter, "A Framework for Large-Scale Measurement of Broadband Performance (LMAP)", [RFC 7594](#), DOI 10.17487/RFC7594, September 2015, <<http://www.rfc-editor.org/info/rfc7594>>.

Authors' Addresses

Joachim Fabini
TU Wien
Gusshausstrasse 25/E389
Vienna 1040
Austria

Phone: +43 1 58801 38813
Fax: +43 1 58801 38898
Email: Joachim.Fabini@tuwien.ac.at
URI: <http://www.tc.tuwien.ac.at/about-us/staff/joachim-fabini/>

Al Morton
AT&T Labs
200 Laurel Avenue South
Middletown, NJ 07748
USA

Phone: +1 732 420 1571
Fax: +1 732 368 1192
Email: acmorton@att.com
URI: <http://home.comcast.net/~acmacm/>