

December 17, 2014

UDP Checksum Trailer in OWAMP and TWAMP
draft-ietf-ippm-checksum-trailer-00.txt

Abstract

The One-Way Active Measurement Protocol (OWAMP) and the Two-Way Active Measurement Protocol (TWAMP) are used for performance monitoring in IP networks. Delay measurement is performed in these protocols by using timestamped test packets. Some implementations use hardware-based timestamping engines that integrate the accurate transmission timestamp into every outgoing OWAMP/TWAMP test packet during transmission. Since these packets are transported over UDP, the UDP checksum field is then updated to reflect this modification. This document proposes to use the last 2 octets of every test packet as a Checksum Trailer, allowing timestamping engines to reflect the checksum modification in the last 2 octets rather than in the UDP checksum field. The behavior defined in this document is completely interoperable with existing OWAMP/TWAMP implementations.

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/ietf/1id-abstracts.txt>.

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

This Internet-Draft will expire on June 17, 2015.

Copyright 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	2
2.	Conventions used in this document	4
2.1.	Terminology	4
2.2.	Abbreviations	4
3.	Using the UDP Checksum Trailer in OWAMP and TWAMP	5
3.1.	Overview	5
3.2.	OWAMP / TWAMP Test Packets with Checksum Trailer	5
3.2.1.	Transmission of OWAMP/TWAMP with Checksum Trailer ..	8
3.2.2.	Intermediate Updates of OWAMP/TWAMP with Checksum Trailer	9
3.2.3.	Reception of OWAMP/TWAMP with Checksum Trailer	9
3.3.	Interoperability with Existing Implementations.....	9
3.4.	Using the Checksum Trailer with or without Authentication	9
4.	Security Considerations	10
5.	IANA Considerations	10
6.	Acknowledgments	10
7.	References	10
7.1.	Normative References	10
7.2.	Informative References	11

[1.](#) Introduction

The One-Way Active Measurement Protocol ([\[OWAMP\]](#)) and the Two-Way Active Measurement Protocol ([\[TWAMP\]](#)) are used for performance monitoring in IP networks.

Delay and delay variation are two of the metrics that OWAMP/TWAMP can measure. This measurement is performed using timestamped test packets.

The accuracy of delay measurements relies on the timestamping method and its implementation. In order to facilitate accurate timestamping, an implementation MAY use a hardware based timestamping engine, as shown in Figure 1. In such cases, the OWAMP/TWAMP packets are sent and received by a software layer, whereas the timestamping engine modifies every outgoing test packet by incorporating its accurate transmission time into the <Timestamp> field in the packet.

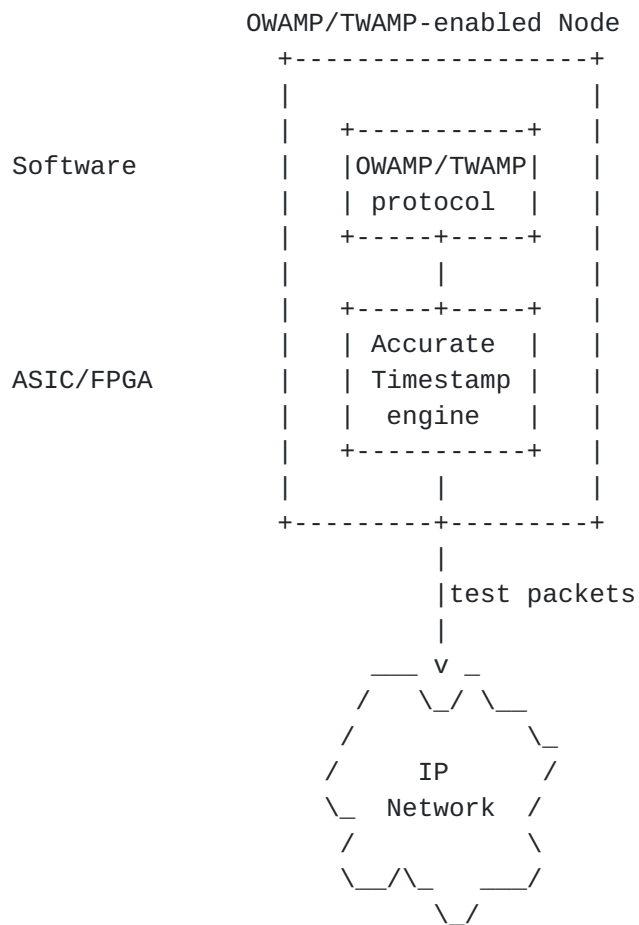


Figure 1 Accurate Timestamping in OWAMP/TWAMP

OWAMP/TWAMP test packets are transported over UDP. When the UDP payload is changed by an intermediate entity such as the timestamping engine, the UDP Checksum field must be updated to reflect the new payload. When using UDP over IPv4 ([[UDP](#)]), an intermediate entity that cannot update the value of the UDP checksum can assign a value of zero to the checksum field, causing the receiver to ignore the

checksum field. UDP over IPv6, as defined in [[IPv6](#)], does not allow a zero checksum, and requires the UDP checksum field to contain a correct checksum of the UDP payload.

Since an intermediate entity only modifies a specific field in the packet, i.e. the timestamp field, the UDP checksum update can be performed incrementally, using the concepts presented in [[Checksum](#)].

A similar problem is addressed in Annex E of [[IEEE1588](#)]. When the Precision Time Protocol (PTP) is transported over IPv6, two octets are appended to the end of the PTP payload for UDP checksum updates. The value of these two octets can be updated by an intermediate entity, causing the value of the UDP checksum field to remain correct.

This document defines a similar concept for [[OWAMP](#)] and [[TWAMP](#)], allowing intermediate entities to update OWAMP/TWAMP test packets and maintain the correctness of the UDP checksum by modifying the last 2 octets of the packet.

The term Checksum Trailer is used throughout this document and refers to the 2 octets at the end of the UDP payload, used for updating the UDP checksum by intermediate entities.

The usage of the Checksum Trailer can in some cases simplify the implementation, since if the packet data is processed in a serial order, it is simpler to first update the timestamp field, and then update the Checksum Trailer rather than to update the timestamp and then update the UDP checksum, residing at the UDP header.

[2. Conventions used in this document](#)

[2.1. 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 [[KEYWORDS](#)].

[2.2. Abbreviations](#)

OWAMP	One-Way Active Measurement Protocol
PTP	Precision Time Protocol
TWAMP	Two-Way Active Measurement Protocol
UDP	User Datagram Protocol

3. Using the UDP Checksum Trailer in OWAMP and TWAMP

3.1. Overview

The UDP Checksum Trailer is a two-octet trailer that is piggybacked at the end of the test packet. It resides in the last 2 octets of the UDP payload.

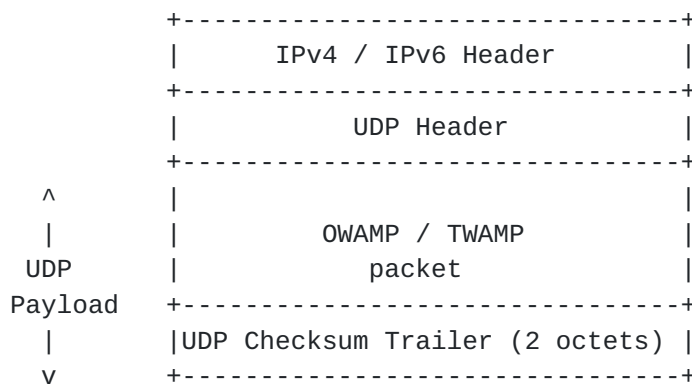


Figure 2 Checksum Trailer in OWAMP/TWAMP Test Packet

3.2. OWAMP / TWAMP Test Packets with Checksum Trailer

The One-Way Active Measurement Protocol [[OWAMP](#)], and the Two-Way Active Measurement Protocol [[TWAMP](#)] both make use of timestamped test packets. The formats of these packets are defined in [[OWAMP](#)] and in [[TWAMP](#)].

OWAMP/TWAMP test packets are transported over UDP, either over IPv4 or over IPv6. This document applies to both OWAMP/TWAMP over IPv4 and over IPv6.

OWAMP/TWAMP test packets contain a Packet Padding field. This document proposes to use the last 2 octets of the Packet Padding field as the Checksum Trailer. In this case the Checksum Trailer is always the last 2 octets of the UDP payload, and thus the trailer is located UDP Length - 2 octets after the beginning of the UDP header.

Figure 3 illustrates the OWAMP test packet format including the UDP checksum trailer.

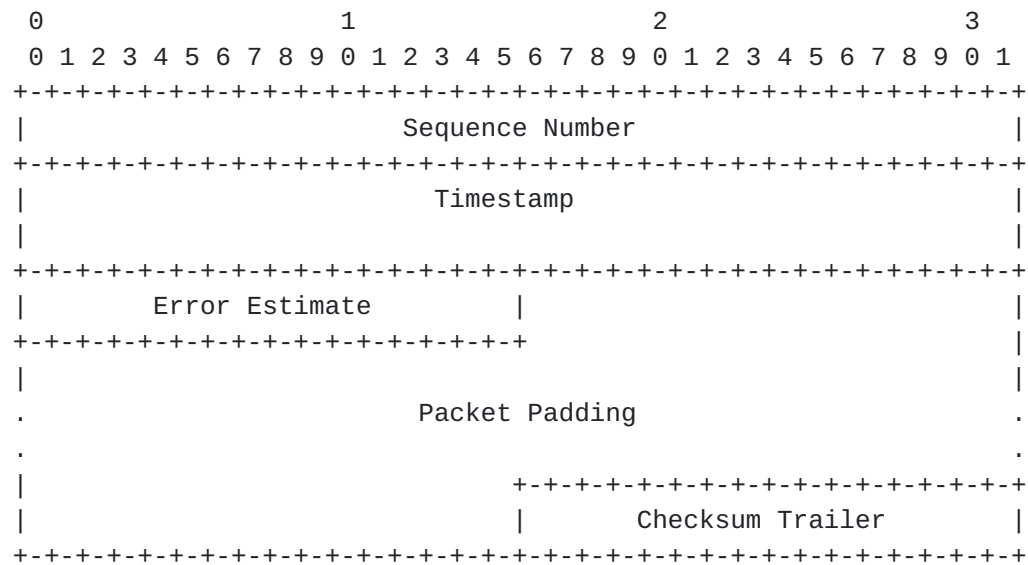


Figure 3 Checksum Trailer in OWAMP Test Packets

Figure 4 illustrates the TWAMP test packet format including the UDP checksum trailer.

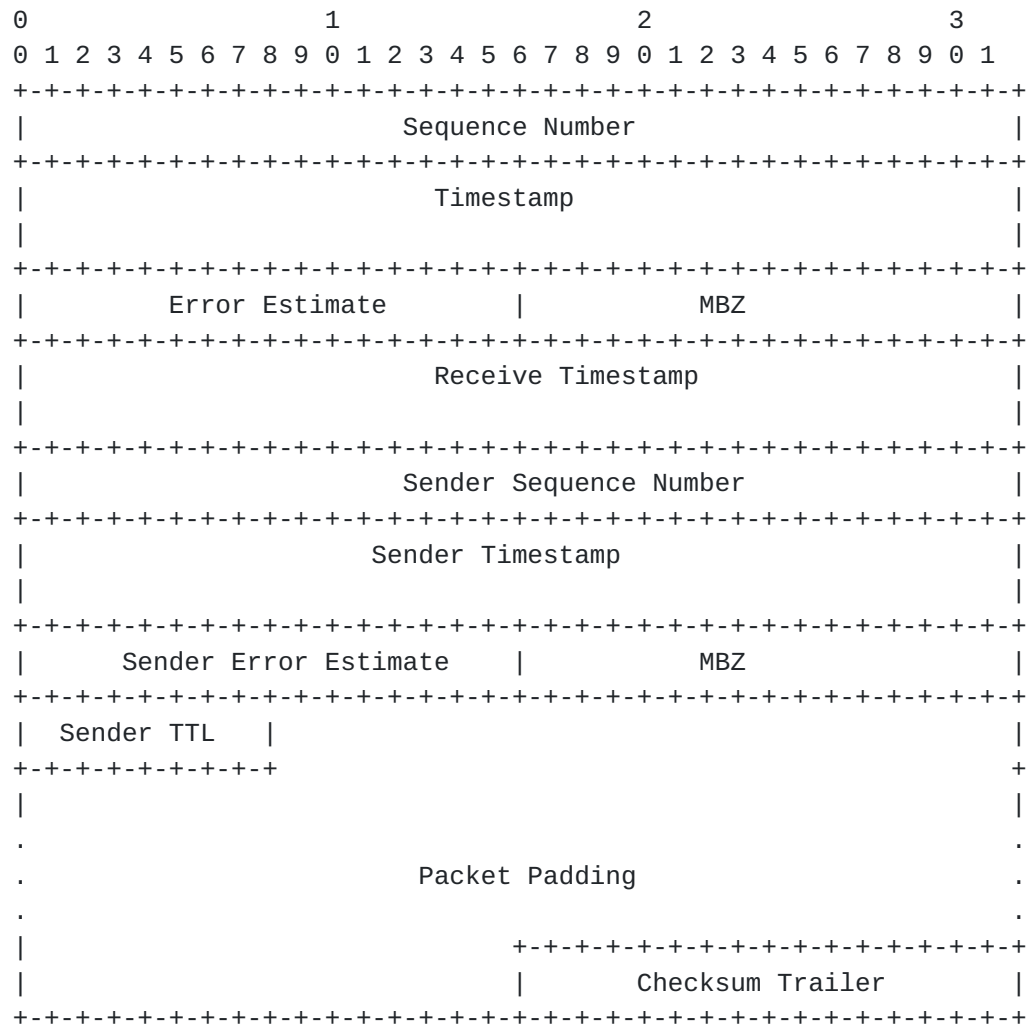


Figure 4 Checksum Trailer in TWAMP Test Packets

The length of the Packet Padding field in test packets is announced during the session initiation through the <Padding Length> field in the Request-Session message [[OWAMP](#)], or in the Request-TW-Session [[TWAMP](#)].

When a Checksum Trailer is included, the <Padding Length> MUST be sufficiently long to include the Checksum Trailer:

- o In OWAMP the padding length is at least 2 octets, allowing the sender to incorporate the checksum trailer in the last 2 octets of the padding.

- o In TWAMP the padding length is at least 29 octets. The additional padding is required since the header of reflector test packets is 27 octets longer than the header of sender test packets. Thus, the padding in reflector test packets is 27 octets shorter than in sender packet. Using 29 octets of padding in sender test packets allows both the sender and the reflector to use a 2-octet checksum trailer.

Note: the 27-octet difference between the sender packet and the reflector packet is specifically in unauthenticated mode, whereas in authenticated mode the difference between the sender and receiver packets is 56 octets. As specified in [Section 3.4.](#) , the checksum trailer should only be used in unauthenticated mode.

- o Two optional TWAMP features are defined in [[RFC6038](#)]: octet reflection and symmetrical size. When at least one of these features is enabled, the Request-TW-Session includes the <Padding Length> field, as well as a <Length of padding to reflect> field. In this case both fields must be sufficiently long to allow at least 2 octets of padding in both sender test packets and reflector test packets.

Specifically, when octet reflection is enabled, the two length fields must be defined such that the padding expands at least 2 octets beyond the end of the reflected octets.

As described in [Section 1.](#) , the extensions described in this document are implemented by two logical layers, a protocol layer and a timestamping layer. It is assumed that the two layers are synchronized about whether the usage of the Checksum Trailer is enabled or not; since both logical layers reside in the same network device, it is assumed there is no need for a protocol that synchronizes this information between the two layers. When Checksum Trailer usage is enabled, the protocol layer must take care to verify that test packets include the necessary padding, and avoiding the need for the timestamping layer to verify that en-route test packets include the necessary padding.

[3.2.1.](#) Transmission of OWAMP/TWAMP with Checksum Trailer

The transmitter of an OWAMP/TWAMP test packet MAY include a Checksum Trailer field, incorporated in the last 2 octets of the Packet Padding.

A transmitter that includes a Checksum Trailer in its outgoing test packets MUST include a Packet Padding in these packets, the length of which MUST be sufficient to include the checksum trailer. The length of the padding field is negotiated during session initiation, as described in [Section 3.2.](#)

3.2.2. Intermediate Updates of OWAMP/TWAMP with Checksum Trailer

An intermediate entity that receives and alters an OWAMP/TWAMP test packet MAY alter the Checksum Trailer field in order to maintain the correctness of the UDP checksum value.

3.2.3. Reception of OWAMP/TWAMP with Checksum Trailer

This document does not impose new requirements on the receiving end of an OWAMP/TWAMP test packet.

The UDP layer at the receiving end verifies the UDP Checksum of received test packets, and the OWAMP/TWAMP layer SHOULD treat the Checksum Trailer as part of the Packet Padding.

3.3. Interoperability with Existing Implementations

The behavior defined in this document does not impose new requirements on the reception behavior of an OWAMP receiver or a TWAMP reflector, since the existence of the checksum trailer is transparent from the perspective of the receiver/reflector. Thus, the functionality described in this document allows interoperability with existing implementations that comply to [\[OWAMP\]](#) or [\[TWAMP\]](#).

3.4. Using the Checksum Trailer with or without Authentication

Both OWAMP and TWAMP may use authentication, as defined in [\[OWAMP\]](#) or [\[TWAMP\]](#). A Checksum Trailer SHOULD NOT be used when authentication is enabled. The Checksum Trailer is effective in unauthenticated mode, allowing the intermediate entity to perform serial processing of the packet without storing-and-forwarding it.

On the other hand, when message authentication is used, an intermediate entity that alters test packets must also re-compute the Message Authentication Code (MAC) accordingly. The MAC update typically requires the intermediate entity to store the packet, re-compute its MAC, and then forward it. Thus, the benefit of the checksum trailer is effectively irrelevant when a MAC is used.

Note: while [\[OWAMP\]](#) and [\[TWAMP\]](#) include an inherent security mechanism, these protocols can be secured by other measures, e.g., [\[IPPMIPsec\]](#). For similar reasons as described above, a Checksum Trailer SHOULD NOT be used in this case.

4. Security Considerations

This document describes how a Checksum Trailer extension can be used for maintaining the correctness of the UDP checksum.

The purpose of this extension is to ease the implementation of accurate timestamping engines, as described in Figure 1. The extension is intended to be used internally in an OWAMP/TWAMP enabled node, and not intended to be used by intermediate switches and routers that reside between the sender and the receiver/reflector. Any modification of a test packet by intermediate switches or routers should be considered a malicious MITM attack.

It is important to emphasize that the scheme described in this document does not increase the protocol's vulnerability to MITM attacks; a MITM who maliciously modifies a packet and its checksum trailer is logically equivalent to a MITM attacker who modifies a packet and its UDP Checksum field.

The concept described in this document is intended to be used only in unauthenticated mode. As described in [Section 3.4.](#), the benefits of the Checksum Trailer do not apply when authentication is enabled.

5. IANA Considerations

There are no IANA actions required by this document.

RFC Editor: please delete this section before publication.

6. Acknowledgments

This document was prepared using 2-Word-v2.0.template.dot.

7. References

7.1. Normative References

- | | |
|------------|--|
| [KEYWORDS] | Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14 , RFC 2119 , March 1997. |
| [IPv6] | Deering, S., Hinden, R., "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460 , December 1998. |
| [Checksum] | Rijsinghani, A., "Computation of the Internet Checksum via Incremental Update", RFC 1624 , May 1994. |

- [UDP] Postel, J., "User Datagram Protocol", [RFC 768](#), August 1980.
- [OWAMP] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and Zekauskas, M., "A One-way Active Measurement Protocol (OWAMP)", [RFC 4656](#), September 2006.
- [TWAMP] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and Babiarz, J., "A Two-Way Active Measurement Protocol (TWAMP)", [RFC 5357](#), October 2008.
- [RFC6038] Morton, A., Ciavattone, L., "Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features", [RFC 6038](#), October 2010.

7.2. Informative References

- [IEEE1588] IEEE TC 9 Instrumentation and Measurement Society, "1588 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems Version 2", IEEE Standard, 2008.
- [IPPMIPsec] Pentikousis, K., Cui, Y., Zhang, E., "Network Performance Measurement for IPsec", [draft-ietf-ippm-ipsec](#) (work in progress), September 2014.

Authors' Addresses

Tal Mizrahi
Marvell
6 Hamada St.
Yokneam, 20692 Israel

Email: talmi@marvell.com