

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 7, 2020

G. Mirsky
X. Min
ZTE Corp.
G. Jun
ZTE Corporation
H. Nydell
Accedian Networks
R. Foote
Nokia
July 6, 2019

Simple Two-way Active Measurement Protocol Optional Extensions
draft-mirsky-ippm-stamp-option-tlv-05

Abstract

This document describes optional extensions to Simple Two-way Active Measurement Protocol (STAMP) which enable measurement performance metrics in addition to ones supported by the STAMP base specification.

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 <https://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 January 7, 2020.

Copyright Notice

Copyright (c) 2019 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 (<https://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	2
2.1. Terminology	2
2.2. Requirements Language	3
3. Theory of Operation	3
4. TLV Extensions to STAMP	4
4.1. Extra Padding TLV	6
4.2. Location TLV	6
4.3. Timestamp Information TLV	8
4.4. Class of Service TLV	9
4.5. Direct Measurement TLV	10
5. IANA Considerations	11
5.1. STAMP TLV Registry	11
5.2. Synchronization Source Sub-registry	12
5.3. Timestamping Method Sub-registry	13
6. Security Considerations	14
7. Acknowledgments	14
8. References	14
8.1. Normative References	14
8.2. Informative References	15
Authors' Addresses	15

1. Introduction

Simple Two-way Active Measurement Protocol (STAMP) [I-D.ietf-ippm-stamp] supports the use of optional extensions that use Type-Length-Value (TLV) encoding. Such extensions are to enhance the STAMP base functions, such as measurement of one-way and round-trip delay, latency, packet loss, as well as ability to detect packet duplication and out-of-order delivery of the test packets. This specification provides definitions of optional STAMP extensions, their formats, and theory of operation.

2. Conventions used in this document

2.1. Terminology

STAMP - Simple Two-way Active Measurement Protocol

DSCP - Differentiated Services Code Point

ECN - Explicit Congestion Notification

NTP - Network Time Protocol

PTP - Precision Time Protocol

HMAC Hashed Message Authentication Code

TLV Type-Length-Value

BITS Building Integrated Timing Supply

SSU Synchronization Supply Unit

GPS Global Positioning System

GLONASS Global Orbiting Navigation Satellite System

LORAN-C Long Range Navigation System Version C

MBZ Must Be Zeroed

CoS Class of Service

2.2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Theory of Operation

STAMP Session-Sender transmits test packets to STAMP Session-Reflector. STAMP Session-Reflector receives Session-Sender's packet and acts according to the configuration and optional control information communicated in the Session-Sender's test packet. STAMP defines two different test packet formats, one for packets transmitted by the STAMP-Session-Sender and one for packets transmitted by the STAMP-Session-Reflector. STAMP supports two modes: unauthenticated and authenticated. Unauthenticated STAMP test packets are compatible on the wire with unauthenticated TWAMP-Test [RFC5357] packet formats.

By default, STAMP uses symmetrical packets, i.e., the size of the packet transmitted by Session-Reflector equals the size of the packet received by the Session-Reflector.

4. TLV Extensions to STAMP

Figure 1 displays the format of STAMP Session-Sender test packet in unauthenticated mode that includes a TLV.




Figure 1: STAMP Session-Sender test packet format with TLV in unauthenticated mode

The MBZ (Must Be Zeroed) field of a test packet transmitted by a STAMP Session-Sender MUST be 30 octets long. A STAMP Session-Sender test packet MUST NOT use the Reflect Octets capability defined in [RFC6038].

TLVs (Type-Length-Value tuples) have the two octets long Type field, two octets long Length field that is the length of the Value field in octets. Type values, see Section 5.1, less than 32768 identify mandatory TLVs that MUST be supported by an implementation. Type values greater than or equal to 32768 identify optional TLVs that SHOULD be ignored if the implementation does not understand or support them. If a Type value for TLV or sub-TLV is in the range for Vendor Private Use, the Length MUST be at least 4, and the first four octets MUST be that vendor's the Structure of Management Information (SMI) Private Enterprise Number, in network octet order. The rest of the Value field is private to the vendor. Following sections

4.1. Extra Padding TLV




Figure 3: Extra Padding TLV

where fields are defined as the following:

- o Extra Padding Type - TBA1 allocated by IANA Section 5.1
- o Length - 2 octets long field equals length on the Extra Padding field in octets.
- o Extra Padding - a pseudo-random sequence of numbers. The field MAY be filled with all zeroes.

The Extra Padding TLV is similar to the Packet Padding field in TWAMP-Test packet [RFC5357]. The in STAMP the Packet Padding field is used to ensure symmetrical size between Session-Sender and Session-Reflector test packets. Extra Padding TLV MUST be used to create STAMP test packets of larger size.

4.2. Location TLV

STAMP session-sender MAY include the Location TLV to request information from the session-reflector. The session-sender SHOULD NOT fill any information fields except for Type and Length. The session-reflector MUST validate the Length value against the address family of the transport encapsulating the STAMP test packet. If the value of the Length field is invalid, the session-reflector MUST zero all fields and MUST NOT return any information to the session-sender. The session-reflector MUST ignore all other fields of the received Location TLV.




Figure 4: Session-Reflector Location TLV

where fields are defined as the following:

- o Location Type - TBA2 allocated by IANA Section 5.1
- o Length - 2 octets long field equals length on the Value field in octets. Length field value MUST be 20 octets for the IPv4 address family. For the IPv6 address family value of the Length field MUST be 44 octets. All other values are invalid.
- o Source MAC - 6 octets 48 bits long field. The session-reflector MUST copy Source MAC of received STAMP packet into this field.
- o Reserved A - two octets long field. MUST be zeroed on transmission and ignored on reception.
- o Destination IP Address - IPv4 or IPv6 destination address of the received by the session-reflector STAMP packet.
- o Source IP Address - IPv4 or IPv6 source address of the received by the session-reflector STAMP packet.
- o Dest.port - one octet long UDP destination port number of the received STAMP packet.
- o Src.port - one octet long UDP source port number of the received STAMP packet.
- o Reserved B - two octets long field. MUST be zeroed on transmission and ignored on reception.

The Location TLV MAY be used to determine the last-hop addressing for STAMP packets including source and destination IP addresses as well as the MAC address of the last-hop router. Last-hop MAC address MAY be monitored by the Session-Sender whether there has been a path switch on the last hop, closest to the Session-Reflector. The IP addresses and UDP port will indicate if there is a NAT router on the path, and allows the Session-Sender to identify the IP address of the Session-Reflector behind the NAT, detect changes in the NAT mapping that could cause sending the STAMP packets to the wrong Session-Reflector.

4.3. Timestamp Information TLV

STAMP session-sender MAY include the Timestamp Information TLV to request information from the session-reflector. The session-sender SHOULD NOT fill any information fields except for Type and Length. The session-reflector MUST validate the Length value of the STAMP test packet. If the value of the Length field is invalid, the session-reflector MUST zero all fields and MUST NOT return any information to the session-sender.




Figure 5: Timestamp Information TLV

where fields are defined as the following:

- o Timestamp Information Type - TBA3 allocated by IANA Section 5.1
- o Length - 2 octets long field, equals four octets.
- o Sync Src In - one octet long field that characterizes the source of clock synchronization at the ingress of Session-Reflector. There are several of methods to synchronize the clock, e.g., Network Time Protocol (NTP) [RFC5905], Precision Time Protocol (PTP) [IEEE.1588.2008], Synchronization Supply Unit (SSU) or Building Integrated Timing Supply (BITS), or Global Positioning System (GPS), Global Orbiting Navigation Satellite System (GLONASS) and Long Range Navigation System Version C (LORAN-C). The value is one of Section 5.2.

- o Timestamp In - one octet long field that characterizes the method by which the ingress of Session-Reflector obtained the timestamp T2. A timestamp may be obtained with hardware assist, via software API from a local wall clock, or from a remote clock (the latter referred to as "control plane"). The value is one of Section 5.3.
- o Sync Src Out - one octet long field that characterizes the source of clock synchronization at the egress of Session-Reflector. The value is one of Section 5.2.
- o Timestamp Out - one octet long field that characterizes the method by which the egress of Session-Reflector obtained the timestamp T3. The value is one of Section 5.3.

4.4. Class of Service TLV

The STAMP session-sender MAY include Class of Service (CoS) TLV in the STAMP test packet. If the CoS TLV is present in the STAMP test packet and the value of the DSCP1 field is zero, then the STAMP session-reflector MUST copy the values of Differentiated Services Code Point (DSCP) ECN fields from the received STAMP test packet into DSCP2 and ECN fields respectively of the CoS TLV of the reflected STAMP test packet. If the value of the DSCP1 field is non-zero, then the STAMP session-reflector MUST use DSCP1 value from the CoS TLV in the received STAMP test packet as DSCP value of STAMP reflected test packet and MUST copy DSCP and ECN values of the received STAMP test packet into DSCP2 and ECN fields of Class of Service TLV in the STAMP reflected a packet. The Session-Sender, upon receiving the reflected packet, will save the DSCP and ECN values for analysis of the CoS in the reverse direction.




Figure 6: Class of Service TLV

where fields are defined as the following:

- o Class of Service Type - TBA4 allocated by IANA Section 5.1
- o Length - 2 octets long field, equals four octets.

- o DSCP1 - The Differentiated Services Code Point (DSCP) intended by the Session-Sender. To be used as the return DSCP from the Session-Reflector.
- o DSCP2 - The received value in the DSCP field at the Session-Reflector in the forward direction.
- o ECN - The received value in the ECN field at the Session-Reflector in the forward direction.
- o Reserved - 18 bits long field, must be zeroed in transmission and ignored on receipt.

A STAMP Session-Sender that includes the CoS TLV sets the value of the DSCP1 field and zeroes the value of the DSCP2 field. A STAMP Session-Reflector that received the test packet with the CoS TLV MUST include the CoS TLV in the reflected test packet. Also, the Session-Reflector MUST copy the value of the DSCP field of the IP header of the received STAMP test packet into the DSCP2 field in the reflected test packet. And, at last, the Session-Reflector MUST set the value of the DSCP field in the IP header of the reflected test packet equal to the value of the DSCP1 field of the test packet it has received.

Re-mapping of CoS in some use cases, for example, in mobile backhaul networks is used to provide multiple services, i.e., 2G, 3G, LTE, over the same network. But if it is misconfigured, then it is often difficult to diagnose the root cause of the problem that is viewed as an excessive packet drop of higher level service while packet drop for lower service packets is at a normal level. Using CoS TLV in STAMP test helps to troubleshoot the existing problem and also verify whether DiffServ policies are processing CoS as required by the configuration.

4.5. Direct Measurement TLV

The Direct Measurement TLV enables collection of "in profile" IP packets that had been transmitted and received by the Session-Sender and Session-Reflector respectfully. The definition of "in-profile packet" is outside the scope of this document.




Figure 7: Direct Measurement TLV

where fields are defined as the following:

- o Direct Measurement Type - TBA5 allocated by IANA Section 5.1
- o Length - 2 octets long field equals length on the Value field in octets. Length field value MUST be 12 octets.
- o Session-Sender Tx counter (S_TxC) is four octets long field.
- o Session-Reflector Rx counter (R_RxC) is four octets long field. MUST be zeroed by the Session-Sender and filled by the Session-Reflector.
- o Session-Reflector Tx counter (R_TxC) is four octets long field. MUST be zeroed by the Session-Sender and filled by the Session-Reflector.

5. IANA Considerations

5.1. STAMP TLV Registry

IANA is requested to create the STAMP TLV Type registry. All code points in the range 1 through 32759 in this registry shall be allocated according to the "IETF Review" procedure as specified in [RFC8126]. Code points in the range 32760 through 65279 in this registry shall be allocated according to the "First Come First Served" procedure as specified in [RFC8126]. Remaining code points are allocated according to Table 1:

Value	Description	Reference
0	Reserved	This document
1- 32767	Mandatory TLV, unassigned	IETF Review
32768 - 65279	Optional TLV, unassigned	First Come First Served
65280 - 65519	Experimental	This document
65520 - 65534	Private Use	This document
65535	Reserved	This document

Table 1: STAMP TLV Type Registry

This document defines the following new values in STAMP TLV Type registry:

Value	Description	Reference
TBA1	Extra Padding	This document
TBA2	Location	This document
TBA3	Timestamp Information	This document
TBA4	Class of Service	This document
TBA5	Direct Measurement	This document

Table 2: STAMP Types

5.2. Synchronization Source Sub-registry

IANA is requested to create Synchronization Source sub-registry as part of STAMP TLV Type registry. All code points in the range 1 through 127 in this registry shall be allocated according to the "IETF Review" procedure as specified in [RFC8126]. Code points in the range 128 through 239 in this registry shall be allocated according to the "First Come First Served" procedure as specified in [RFC8126]. Remaining code points are allocated according to Table 1:

Value	Description	Reference
0	Reserved	This document
1- 127	Unassigned	IETF Review
128 - 239	Unassigned	First Come First Served
240 - 249	Experimental	This document
250 - 254	Private Use	This document
255	Reserved	This document

Table 3: Synchronization Source Sub-registry

This document defines the following new values in Synchronization Source sub-registry:

Value	Description	Reference
1	NTP	This document
2	PTP	This document
3	SSU/BITS	This document
4	GPS/GLONASS/LORAN-C	This document
5	Local free-running	This document

Table 4: Synchronization Sources

5.3. Timestamping Method Sub-registry

IANA is requested to create Timestamping Method sub-registry as part of STAMP TLV Type registry. All code points in the range 1 through 127 in this registry shall be allocated according to the "IETF Review" procedure as specified in [RFC8126]. Code points in the range 128 through 239 in this registry shall be allocated according to the "First Come First Served" procedure as specified in [RFC8126]. Remaining code points are allocated according to Table 1:

Value	Description	Reference
0	Reserved	This document
1- 127	Unassigned	IETF Review
128 - 239	Unassigned	First Come First Served
240 - 249	Experimental	This document
250 - 254	Private Use	This document
255	Reserved	This document

Table 5: Timestamping Method Sub-registry

This document defines the following new values in Timestamping Methods sub-registry:

Value	Description	Reference
1	HW assist	This document
2	SW local	This document
3	Control plane	This document

Table 6: Timestamping Methods

6. Security Considerations

Use of HMAC in authenticated mode may be used to simultaneously verify both the data integrity and the authentication of the STAMP test packets.

7. Acknowledgments

Authors much appreciate the thorough review and thoughtful comments received from Tianran Zhou.

8. References

8.1. Normative References

[I-D.ietf-ippm-stamp]

Mirsky, G., Jun, G., Nydell, H., and R. Foote, "Simple Two-way Active Measurement Protocol", draft-ietf-ippm-stamp-06 (work in progress), April 2019.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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, <<https://www.rfc-editor.org/info/rfc5357>>.
- [RFC6038] Morton, A. and L. Ciavattone, "Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features", RFC 6038, DOI 10.17487/RFC6038, October 2010, <<https://www.rfc-editor.org/info/rfc6038>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

8.2. Informative References

- [IEEE.1588.2008] "Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", IEEE Standard 1588, March 2008.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/info/rfc5905>>.

Authors' Addresses

Greg Mirsky
ZTE Corp.

Email: gregimirsky@gmail.com

Xiao Min
ZTE Corp.

Email: xiao.min2@zte.com.cn

Guo Jun
ZTE Corporation
68# Zijinghua Road
Nanjing, Jiangsu 210012
P.R.China

Phone: +86 18105183663
Email: guo.jun2@zte.com.cn

Henrik Nydell
Accedian Networks

Email: hnydell@accedian.com

Richard Foote
Nokia

Email: footer.foote@nokia.com