

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 2, 2018

G. Mirsky
ZTE Corp.
G. Jun
ZTE Corporation
H. Nydell
Accedian Networks
October 29, 2017

Simple Two-way Active Measurement Protocol
draft-mirsky-ippm-stamp-01

Abstract

This document describes a Two-way Active Measurement Protocol which enables measurement of both one-way and round-trip performance metrics like delay, delay variation and packet loss.

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 May 2, 2018.

Copyright Notice

Copyright (c) 2017 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. Softwarization of Performance Measurement	3
4. Theory of Operation	3
4.1. Sender Behavior and Packet Format	4
4.1.1. Sender Packet Format in Unauthenticated Mode	4
4.1.2. Sender Packet Format in Authenticated and Encrypted Modes	6
4.2. Reflector Behavior and Packet Format	7
4.2.1. Reflector Packet Format in Unauthenticated Mode	7
4.2.2. Reflector Packet Format in Authenticated and Encrypted Modes	8
5. TLV Extensions to STAMP	9
5.1. Extra Padding TLV	10
5.2. Location TLV	10
5.3. Timestamp Information TLV	12
5.4. Class of Service TLV	12
6. IANA Considerations	13
6.1. STAMP TLV Registry	13
6.2. Synchronization Source Sub-registry	14
6.3. Timestamp Method Sub-registry	14
6.4. CoS Operation Sub-registry	14
7. Security Considerations	14
8. Acknowledgments	14
9. Normative References	14
Authors' Addresses	15

1. Introduction

2. Conventions used in this document

2.1. Terminology

STAMP - Simple Two-way Active Measurement Protocol

NTP - Network Time Protocol

PTP - Precision Time Protocol

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. Softwarization of Performance Measurement

Instance of a Simple Two-way Active Measurement Protocol (STAMP) session between a Sender and a Reflector controlled by communication between a Configuration Client as a manager and Configuration Servers as agents of the configuration session that configures STAMP measurement between Sender and Reflector. The Configuration Client also issues queries to obtain operational state information and/or measurement results.

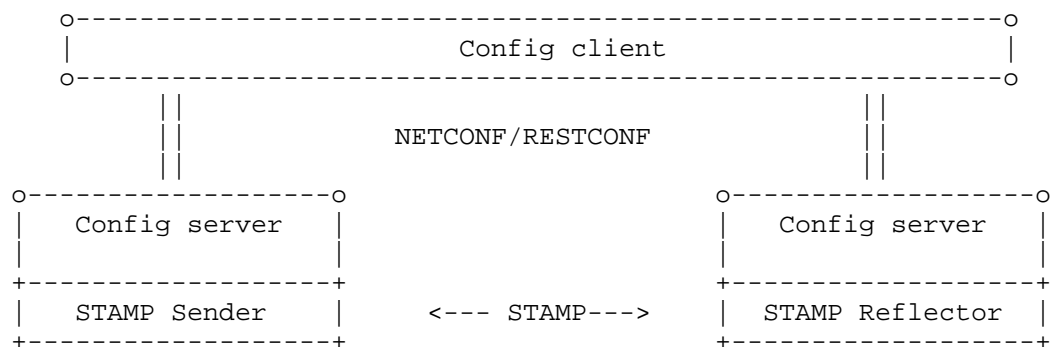


Figure 1: STAMP Reference Model

4. Theory of Operation

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

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

4.1. Sender Behavior and Packet Format

4.1.1. Sender Packet Format in Unauthenticated Mode

Because STAMP supports symmetrical test packets, STAMP Sender packet has minimum size of 44 octets in unauthenticated mode, see Figure 2, and 48 octets in authenticated or encrypted modes , see Figure 4.

For unauthenticated mode:

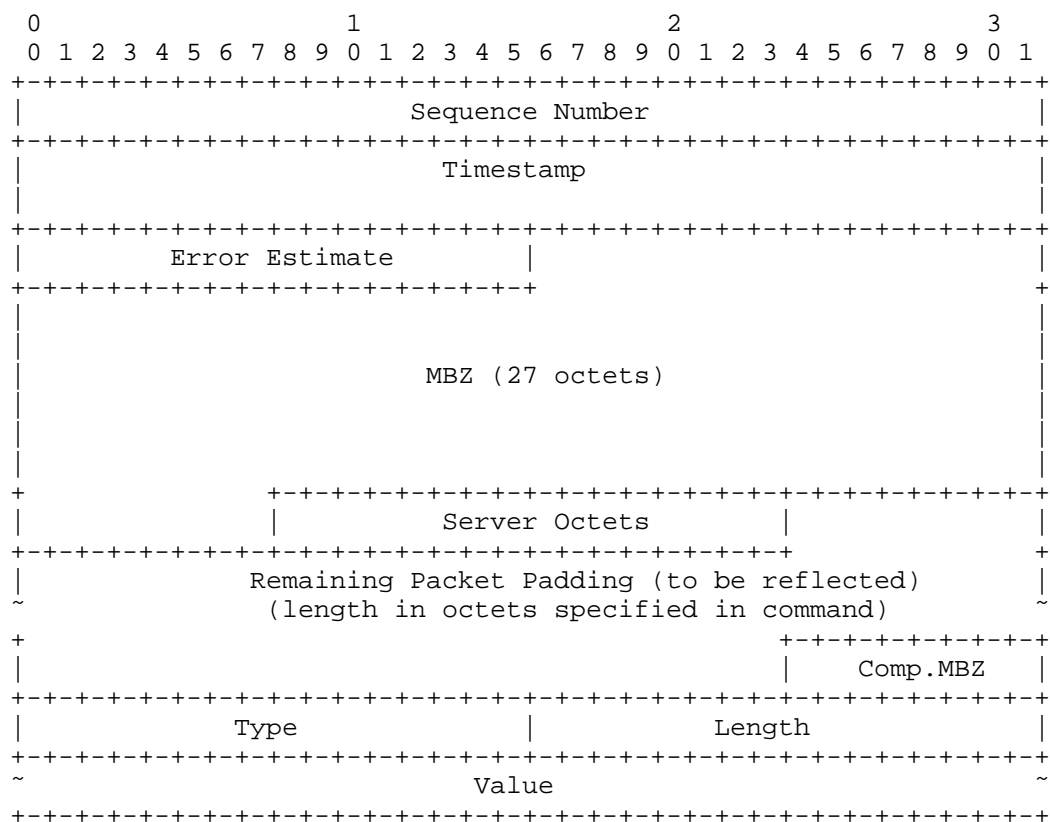


Figure 2: STAMP Sender test packet format in unauthenticated mode

where fields are defined as the following:

- o Sequence Number is four octets long field. For each new session its value starts at zero and is incremented with each transmitted packet.
- o Timestamp is eight octets long field. STAMP node MUST support Network Time Protocol (NTP) version 4 64-bit timestamp format [RFC5905]. STAMP node MAY support IEEE 1588v2 Precision Time Protocol truncated 64-bit timestamp format [IEEE.1588.2008].
- o Error Estimate is two octets long field with format displayed in Figure 3




Figure 3: Error Estimate Format

where S, Scale and Multiplier fields are interpreted as they have been defined in section 4.1.2 [RFC4656]; and Z field - as has been defined in section 2.3 [RFC8186]:

- * 0 - NTP 64 bit format of a timestamp;
- * 1 - PTPv2 truncated format of a timestamp.

- o Must-be-Zero (MBZ) field in the sender unauthenticated packet is 27 octets long. It MUST be all zeroed on transmission and ignored on receipt.
- o Server Octets field is two octets long field. It MUST follow the 27 octets long MBZ field. The Reflect Octets capability defined in [RFC6038]. The value in the Server Octets field equals to the number of octets the Reflector is expected to copy back to the Sender starting with the Server Octets field. Thus the minimal non-zero value for the Server Octets field is two and value of one is invalid. If none of Payload to be copied the value of the Server Octets field MUST be set to zero on transmit.
- o Remaining Packet Padding is optional field of variable length. The number of octets in the Remaining Packet Padding field is the value of the Server Octets field less the length of the Server Octets field.
- o Comp.MBZ is variable length field used to achieve alignment on word boundary. Thus the length of Comp.MBZ field may be only 0,

1, 2 or 3 octets. The value of the field MUST be zeroed on transmission and ignored on receipt.

The unauthenticated STAMP Sender packet MAY include Type-Length-Value encodings that immediately follow the Comp. MBZ field.

- o Type field is two octets long. The value of the Type field is the codepoint allocated by IANA Section 6 that identifies data in the Value field.
- o Length is two octets long field and its value is the length of the Value field in octets.

4.1.2. Sender Packet Format in Authenticated and Encrypted Modes

For authenticated and encrypted modes:

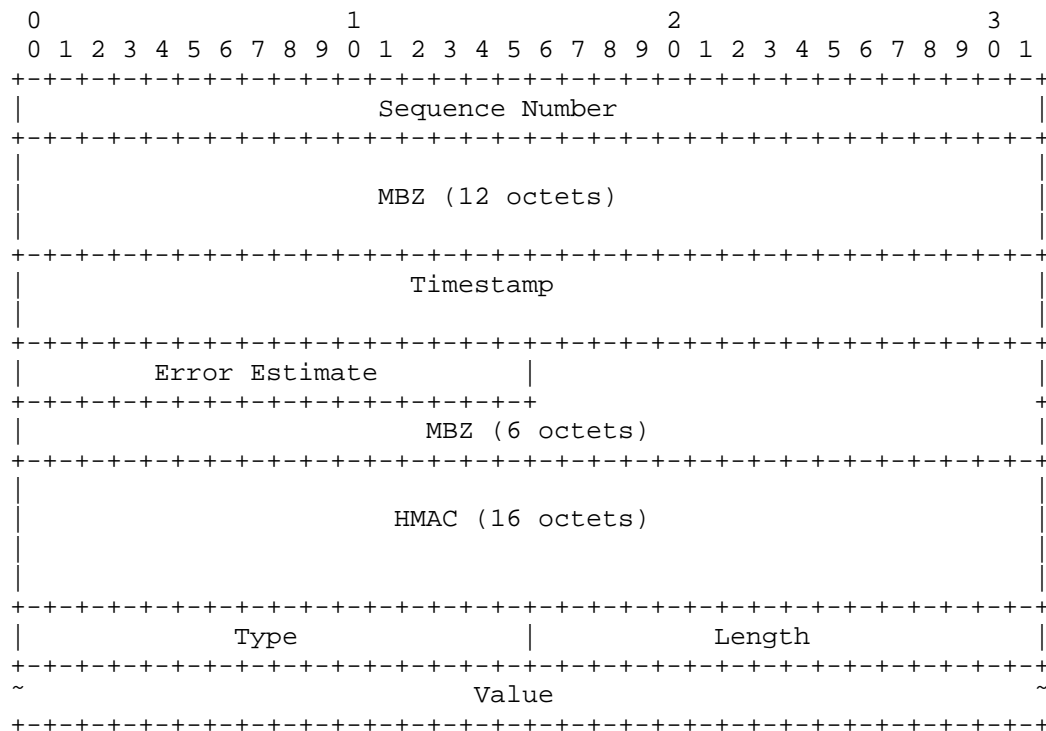


Figure 4: STAMP Sender test packet format in authenticated or encrypted modes

4.2. Reflector Behavior and Packet Format

The Reflector receives the STAMP test packet, verifies it, prepares and transmits the reflected test packet. [Editor note: Verification may include presence and content of TLVs in the STAMP test packet.]

4.2.1. Reflector Packet Format in Unauthenticated Mode

For unauthenticated mode:

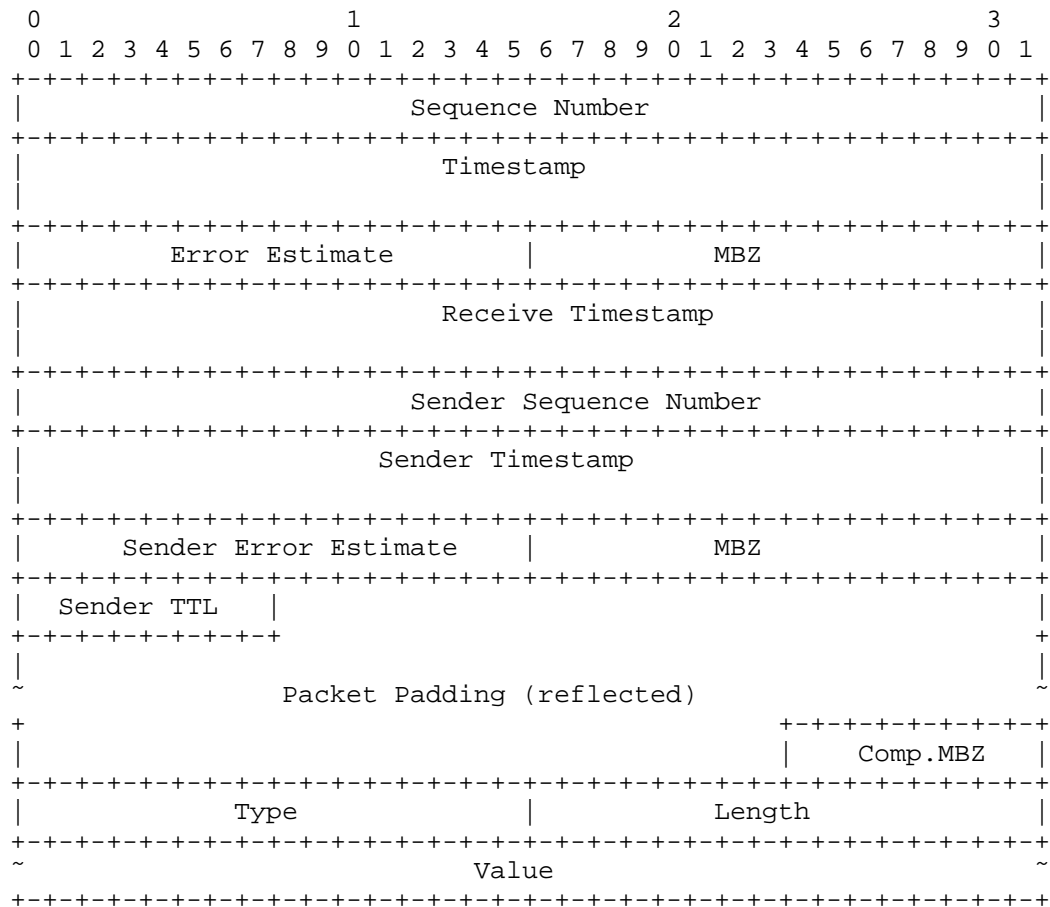



Figure 5: STAMP Reflector test packet format in unauthenticated mode where fields are defined as the following:

- o Sequence Number is four octets long field. The value of the Sequence Number field is set according to the mode of the STAMP Reflector:
 - * in the stateless mode the Reflector copies the value from the received STAMP test packet's Sequence Number field;
 - * in the stateful mode the Reflector counts the received STAMP test packets in each test session and uses that counter to set value of the Sequence Number field.
- o Timestamp and Receiver Timestamp fields are each 8 octets long. The format of these fields, NTP or PTPv2, indicated by the Z flag of the Error Estimate field as described in Section 4.1.
- o Error Estimate has the same size and interpretation as described in Section 4.1.
- o Sender Sequence Number, Sender Timestamp, and Sender Error Estimate are copies of the corresponding fields in the STAMP test packet send by the Sender.
- o Sender TTL is one octet long field and its value is the copy of the TTL field from the received STAMP test packet.
- o Packet Padding (reflected) is optional variable length field. The length of the Packet Padding (reflected) field MUST be equal to the value of the Server Octets field (Figure 2). If the value is non-zero, the Reflector copies octets starting with the Server Octets field.
- o Comp.MBZ is variable length field used to achieve alignment on word boundary. Thus the length of Comp.MBZ field may be only 0, 1, 2 or 3 octets. The value of the field MUST be zeroed on transmission and ignored on receipt.

4.2.2. Reflector Packet Format in Authenticated and Encrypted Modes

For authenticated and encrypted modes:



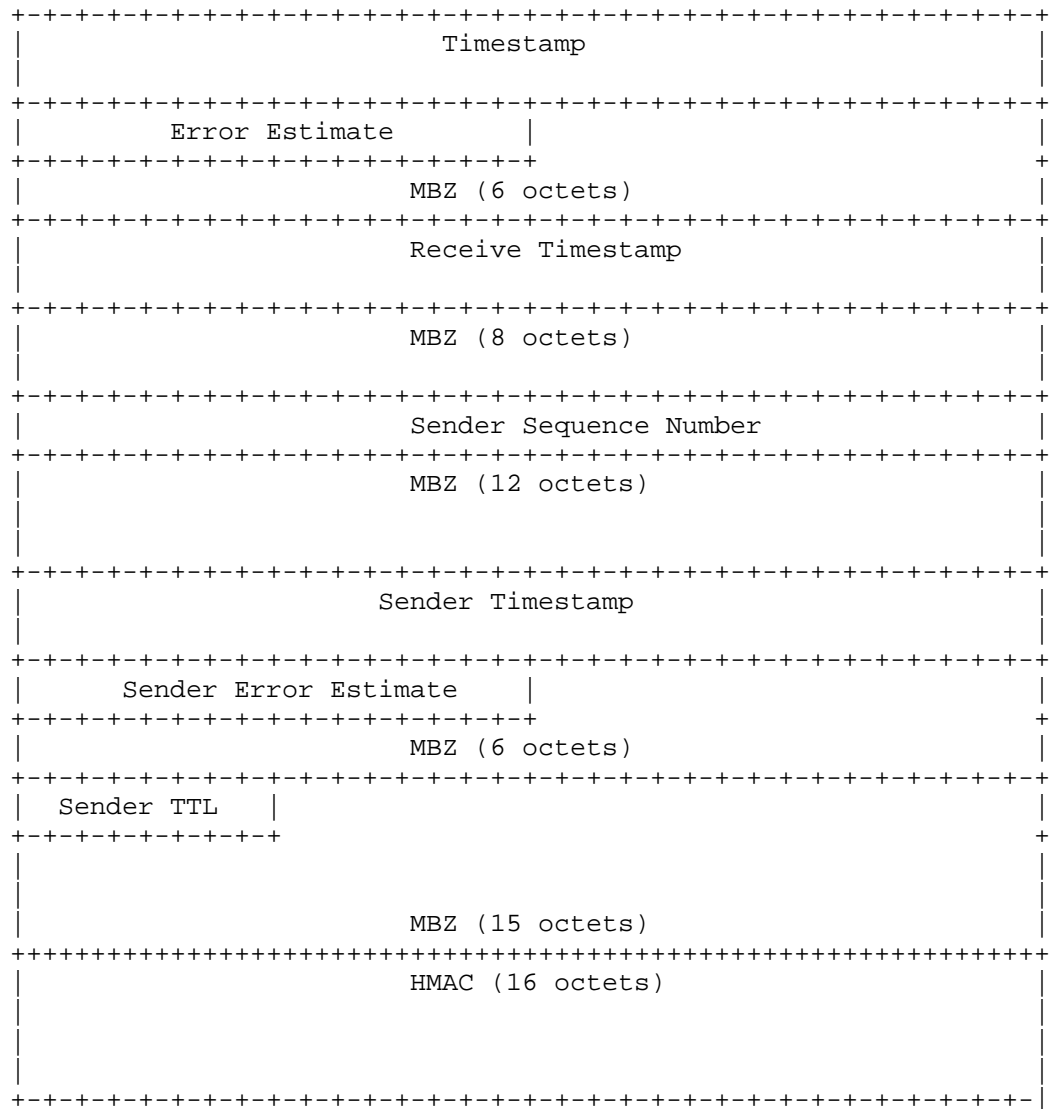


Figure 6: STAMP Reflector test packet format in authenticated or encrypted modes

5. TLV Extensions to STAMP

TBA

5.1. Extra Padding TLV

TBA




Figure 7: Extra Padding TLV

where fields are defined as the following:

- o Extra Padding Type - TBA1 allocated by IANA Section 6.1
- o Length - 2 octets long field equals length on the Extra Padding field in octets.
- o Extra Padding - pseudo-random sequence of numbers. The field MAY be filled with all zeroes.

5.2. Location TLV

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

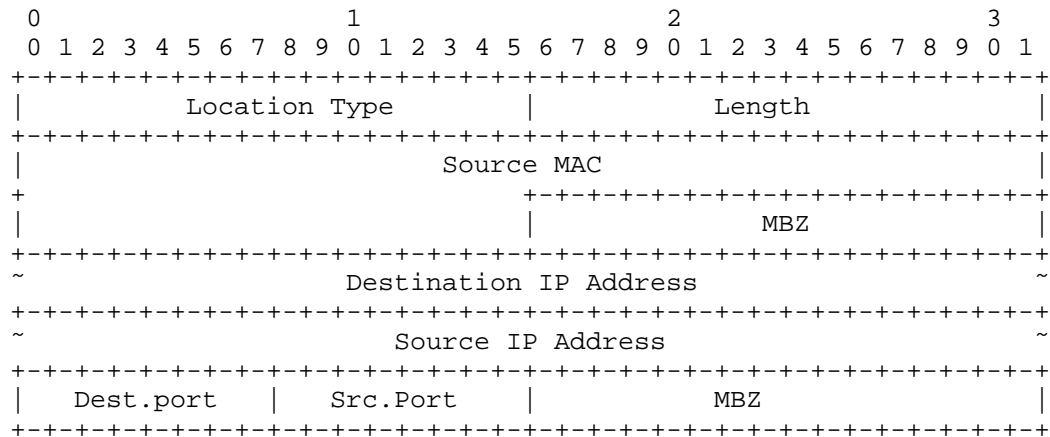


Figure 8: Reflector Location TLV

where fields are defined as the following:

- o Location Type - TBA1 allocated by IANA Section 6.1
- o Length - 2 octets long field equals length on the Value field in octets. Length field value MUST be 20 octets for IPv4 address family. For 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 reflector MUST copy Source MAC of received STAMP packet into this field.
- o MBZ - 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 reflector STAMP packet.
- o Source IP Address - IPv4 or IPv6 source address of the received by the 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.

5.3. Timestamp Information TLV

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




Figure 9: Timestamp Information TLV

where fields are defined as the following:

- o Timestamp Information Type - TBA1 allocated by IANA Section 6.1
- o Length - 2 octets long field, equals 4 octets.
- o Synchronization Source - two octets long field that characterizes the source of clock synchronization at the reflector. The value is one of Section 6.2.
- o Timestamp Method - two octets long field that characterizes timestamping method at the reflector. The value is one of Section 6.3. [Ed.note: Should it be split for ingress and egress?]

5.4. Class of Service TLV

The STAMP sender MAY include Class of Service TLV in the STAMP test packet. If the Class of Service TLV is present in the STAMP test packet and the value of the Op field equals Report (TBA5) value Section 6.4, then the STAMP reflector MUST copy DSCP and ECN values from the received STAMP test packet into DSCP and ECN fields of the Class of Service TLV of the reflected STAMP test packet. If the value of the Op field equals Set and Report (TBA6) Section 6.4, then the STAMP reflector MUST use DSCP value from the Class of Service 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 DSCP and ECN fields of Class of Service TLV in the STAMP reflected packet.




Figure 10: Class of Service TLV

where fields are defined as the following:

o

6. IANA Considerations

6.1. STAMP TLV Registry

IANA is requested to create 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 the Table 1:

Value	Description	Reference
0	Reserved	This document
1- 32759	Unassigned	IETF Review
32760 - 65279	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

Table 2: STAMP Types

6.2. Synchronization Source Sub-registry

TBD

6.3. Timestamp Method Sub-registry

TBD

6.4. CoS Operation Sub-registry

TBD

7. Security Considerations

TBD

8. Acknowledgments

TBD

9. Normative References

- [IEEE.1588.2008]
"Standard for a Precision Clock Synchronization Protocol
for Networked Measurement and Control Systems",
IEEE Standard 1588, March 2008.
- [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>>.
- [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,
<<https://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, <<https://www.rfc-editor.org/info/rfc5357>>.
- [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>>.
- [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>>.
- [RFC8186] Mirsky, G. and I. Meilik, "Support of the IEEE 1588 Timestamp Format in a Two-Way Active Measurement Protocol (TWAMP)", RFC 8186, DOI 10.17487/RFC8186, June 2017, <<https://www.rfc-editor.org/info/rfc8186>>.

Authors' Addresses

Greg Mirsky
ZTE Corp.

Email: gregimirsky@gmail.com

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

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

Internet-Draft

STAMP

October 2017

Henrik Nydell
Accedian Networks

Email: hnydell@accedian.com