Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 6, 2021

Z. Li
China Mobile
M. Chen
Huawei
G. Mirsky
ZTE Corp.

November 02, 2020

Performance Measurement on LAG
draft-li-ippm-pm-on-lag-02

Abstract

This document defines extensions to One-way Active Measurement Protocol (OWAMP), Two-way Active Measurement Protocol (TWAMP), and Simple Two-Way Active Measurement Protocol (STAMP) to implement performance measurement on every member link of a Link Aggregation Group (LAG). With the measured metrics of each member links of a LAG, it enables operators to enforce performance metric based traffic steering policy among the member links.

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 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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 6, 2021.

Table of Contents

1. Problem Statement

Link Aggregation Group (LAG), as defined in [IEEE802.1AX], provides mechanisms to combine multiple physical links into a single logical link. This logical link provides higher bandwidth and better resiliency, because if one of the physical member links fails, the

aggregate logical link can continue to forward traffic over the remaining operational physical member links.

Normally, when forwarding traffic over a LAG, a hash based or the like mechanism is used to load balance the traffic among member links of the LAG. In some cases, the link delays of the member links are different because the member links are over different transport paths. To provide low delay service to time sensitive traffic, we have to know the link delay of each member link of a LAG and then steer traffic accordingly. This requires a solution that could measure the performance metrics of each member link of a LAG.

However, when using One-way Active Measurement Protocol (OWAMP) [RFC4656], Two-way Active Measurement Protocol (TWAMP) [RFC5357], or Simple Two-Way Active Measurement Protocol (STAMP) [RFC8762] to measure the performance of a LAG, the LAG is treated as a single logical link/path. The measured metrics reflect the performance of one member link or an average of some/all member links of the LAG.

In addition, for LAG, using passive or hybrid methods (like alternative marking[RFC8321] or iOAM [I-D.ietf-ippm-ioam-data]) can only monitor the link crossed by traffic. Means the measured metrics only reflect the performance of some member links or an average of some/all member links of the LAG as well. Therefore, in order to measure every link of a LAG, using active methods would be more appropriate.

This document defines extensions to OWAMP [RFC4656], TWAMP [RFC5357] or STAMP [RFC8762] to implement performance measurement on every member link of a LAG.

2. Micro Session on LAG

This document intends to address the scenario (e.g., Figure 1) where two hosts (A and B) are directly connected by a LAG (e.g., the LAG is consisted by three links). The purpose is to measure the performance

of each link of the LAG.

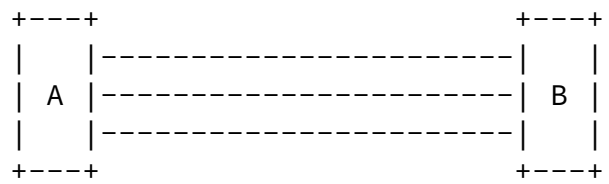

Figure 1: PM for LAG

To measure performance metrics of every member link of a LAG, multiple sessions (one session for each member link) need to be

established between the two hosts that are connected by the LAG. These sessions are called micro sessions in the remainder of this document.

All micro sessions of a LAG share the same Sender Address, Receiver Address. As for the Sender Port and Receiver Port, the micro sessions may share the same Sender Port and Receiver Port pair, or each micro session is configured with different Sender Port and Receiver Port pair. But from simplifying operation point of view, the former is recommended.

In addition, with micro sessions, there needs a way to correlate a session with a member link. For example, when receives a Control or Test packet, the Server/Reflector/Receiver needs to know from which member link the packet is received, and then correlate the packet with a micro session. This is different from the existing OWAMP [RFC4656], TWAMP [RFC5357], or STAMP [RFC8762].

This document defines new command types to indicate that a session is a micro session, the details are described in Section 3 and 4 of this document. For a micro session, on receiving of a Control/Test packet, the receiver uses the receiving link to correlate the packet with a particular session. In addition, Test packets may need to carry the member link information for validation checking. For example, when a Session-Sender receives a Test packet, it may need to check whether the Test packet is from the expected member link.

3. Mirco OWAMP Session

This document assumes that the OWAMP Server and the OWAMP Receiver of an OWAMP micro session are at the same host.

3.1. Micro OWAMP-Control

To support micro OWAMP session, a new command, which is referred to as Request-OW-Micro-Session (TBD1), is defined in this document. The Request-OW-Micro-Session command is based on the OWAMP Request-Session command, and uses the message format as described in Section 3.5 of OWAMP [RFC4656]. Test session creation of micro OWAMP session follows the same procedure as defined in Section 3.5 of OWAMP [RFC4656] with the following additions:

When a OWAMP Server receives a Request-OW-Micro-Session command, if the Session is accepted, the OWAMP Server MUST build an association between the session and the member link from which the Request-Session message is received.

3.2. Micro OWAMP-Test

Micro OWAMP-Test reuses the OWAMP-Test packet format and procedures as defined in Section 4 of OWAMP [RFC4656] with the following additions:

The micro OWAMP Sender MUST send the micro OWAMP-Test packets over the member link with which the session is associated. When receives a Test packet, the micro OWAMP receiver MUST use the member link from which the Test packet is received to correlate the micro OWAMP session. If there is no such a session, the Test packet MUST be discarded.

4. Mirco TWAMP Session

As above, this document assumes that the TWAMP Server and the TWAMP Session-Reflector of a micro OWAMP session are at the same host.

4.1. Micro TWAMP-Control

To support micro TWAMP session, a new command, which is referred to as Request-TW-Micro-Session (TBD2), is defined in this document. The

Request-TW-Micro-Session command is based on the TWAMP Request-Session command, and uses the message format as described in Section 3.5 of TWAMP [RFC5357]. Test session creation of micro TWAMP session follows the same procedure as defined in Section 3.5 of TWAMP [RFC5357] with the following additions:

When a micro TWAMP Server receives a Request-TW-Micro-Session command, if the micro TWAMP Session is accepted, the micro TWAMP Server MUST build an association between the session and the member link from which the Request-Session message is received.

4.2. Micro TWAMP-Test

The micro TWAMP-Test protocol is based on the TWAMP-Test protocol [RFC5357] with the following extensions.

4.2.1. Sender Behavior

In addition to inheriting the TWAMP sender behavior as defined Section 4.1 of [RFC5357], the micro TWAMP Session-Sender MUST send the micro TWAMP-Test packets over the member link with which the session is associated.

When sending Test packet, the micro TWAMP Session-Sender MUST put the Sender member link identifier that is associated with the micro TWAMP session in the Sender Member Link ID. If the Session-Sender knows

the Reflector member link identifier, it MUST put it in the Reflector Member Link ID fields (see Figure 2 and Figure 3). Otherwise, the Reflector Member Link ID field MUST be set to zero.

The Sender member link identifier is used by the Session-Sender to check whether a reflected Test packet is received from the member link that associates to the correct micro TWAMP session. The Reflector member link identifier is used by the Session-Receiver to check whether a Test packet is received from the member link that associates to the correct micro TWAMP session.

The Reflector member link identifier can be obtained from pre-configuration or learned through control plane or data plane (e.g., learned from a reflected Test packet). How to abtain/learn the Reflector member link identifier is out of the scope of this

document.

When receives a reflected Test packet, the micro TWAMP Session-Sender MUST use the member link from which the Test packet is received to correlate to a micro TWAMP session and use the Sender member link identifier to validate whether the Test packet is correctly transmitted over the expected member link. If there is no such a micro TWAMP session, or the validation is failed, the Test packet MUST be discarded.

4.2.1.1. Packet Format and Content

The micro TWAMP Session-Sender packet format is based on the TWAMP Session-Sender packet format as defined in Section 4.1.2 of [RFC5357]. In addition, in order to carry the LAG member link identifier, two new fields (Sender and Reflector Member Link ID) are added. The formats are as below:

For unauthenticated mode:

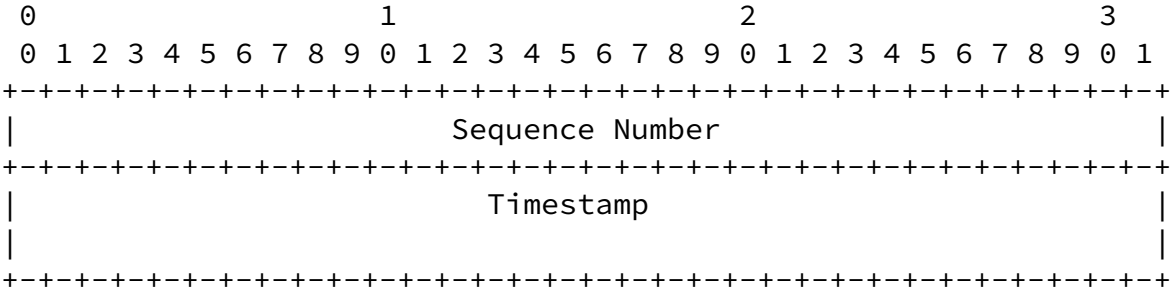

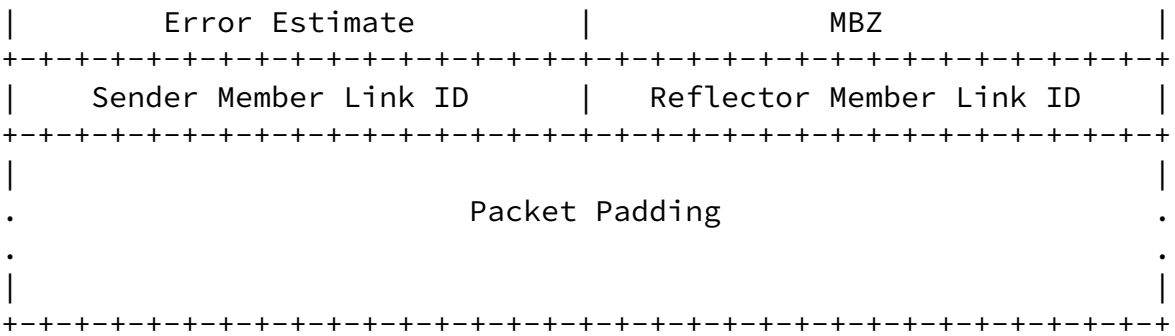

Figure 2: Session-Sender Packet format in Unauthenticated Mode

For authenticated mode:

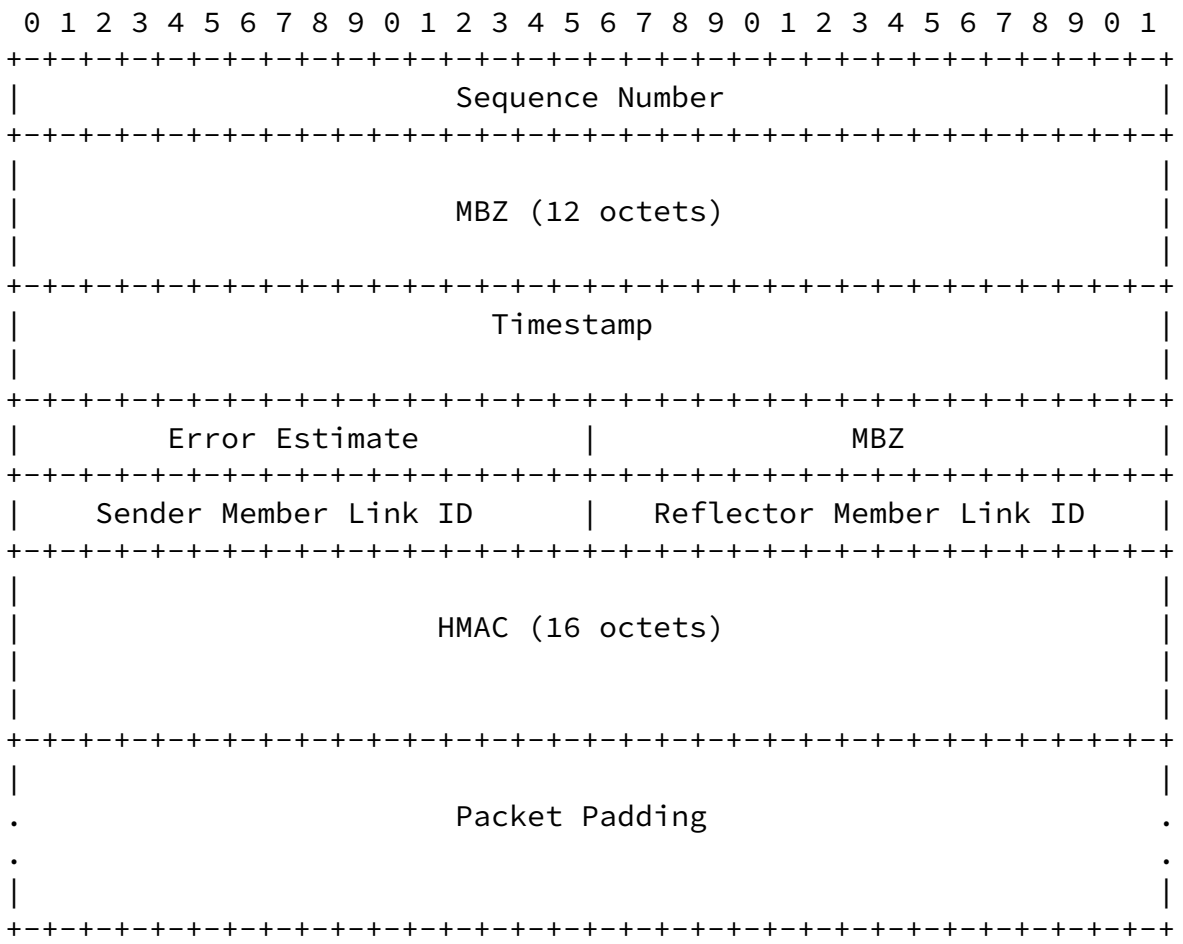

Figure 3: Session-Sender Packet Format in Authenticated Mode

Except for the Sender/Reflector Member Link ID field, all the other fields are the same as defined in Section 4.1.2 of TWAMP [RFC5357], which is originally defined in Section 4.1.2 of OWAMP [RFC4656]. Therefore, it follows the same procedure and guidelines as defined in Section 4.1.2 of TWAMP [RFC5357].

Sender Member Link ID (2-octets in length): it is defined to carry the LAG member link identifier of the Sender side. The value of the Sender Member Link ID MUST be unique at the Session-Sender.

Reflector Member Link ID (2-octets in length): it is defined to carry the LAG member link identifier of the Reflector side. The value of the Reflector Member ID MUST be unique at the Session-Reflector.

4.2.2. Reflector Behavior

The micro TWAMP Session-Reflector inherits the behaviors of a TWAMP Session-Reflector as defined in Section 4.2 of [RFC5357].

In addition, when receives a Test packet, the micro TWAMP Session-Reflector MUST use the member link from which the Test packet is received to correlate to a micro TWAMP session. If there is no such a session, the Test packet MUST be discarded. If Reflector Member Link ID is not zero, the Reflector MUST use the Reflector member link identifier to check whether it associates with the member link from which the Test packet is received. If no, the Test packet MUST be discarded.

When sends a response to the received Test packet, the micro TWAMP Session-Sender MUST copy the Sender member link identifier from the received Test packet and put it in the Sender Member Link ID field of the reflected Test packet (see Figure 4 and Figure 5). In addition, the micro TWAMP Session-Sender MUST put the Reflector member link identifier that are associated with the micro TWAMP session in the and Reflector Member Link ID fields (see Figure 4 and Figure 5).

4.2.2.1. Packet Format and Content

The micro TWAMP Session-Reflector packet format is based on the TWAMP Session-Reflector packet format as defined in Section 4.2.1 of [RFC5357]. In addition, in order to carry the LAG member link identifier, two new fields (Sender and Reflector Member Link ID) are added. The formats are as below:

For unauthenticated mode:

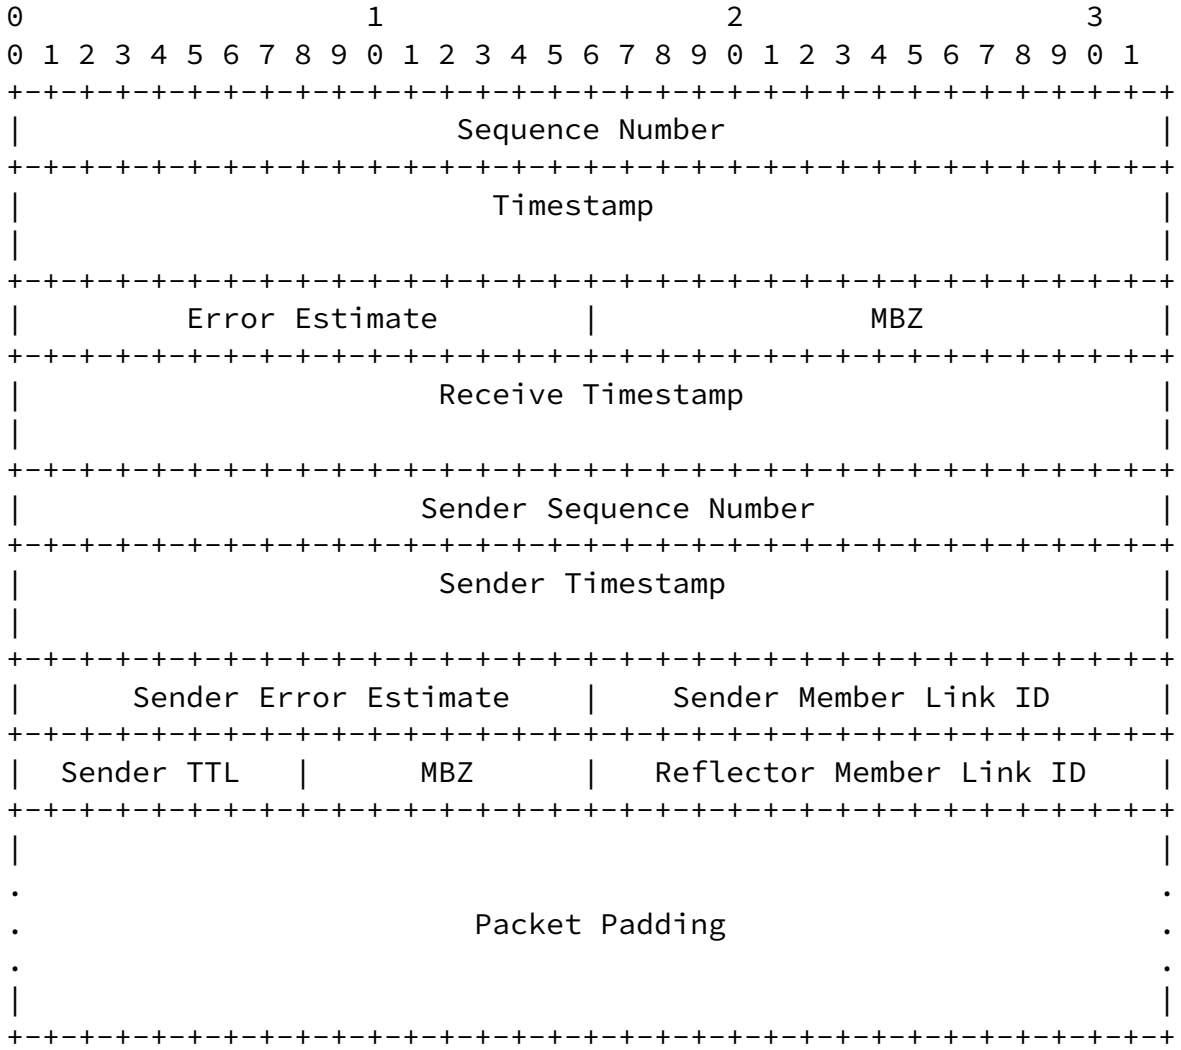

Figure 4: Session-Reflector Packet Format in Unauthenticated Mode

For authenticated and encrypted modes:

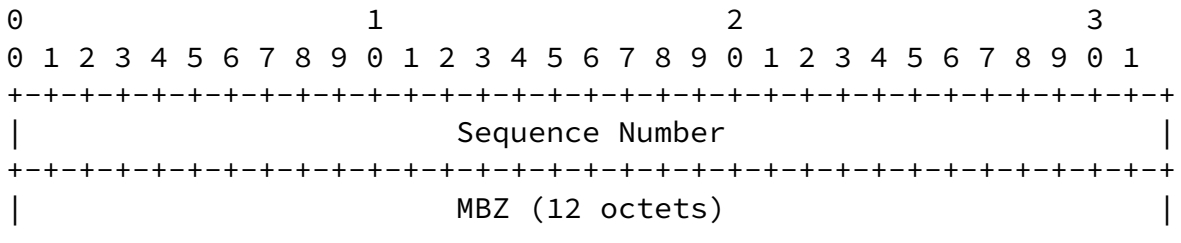

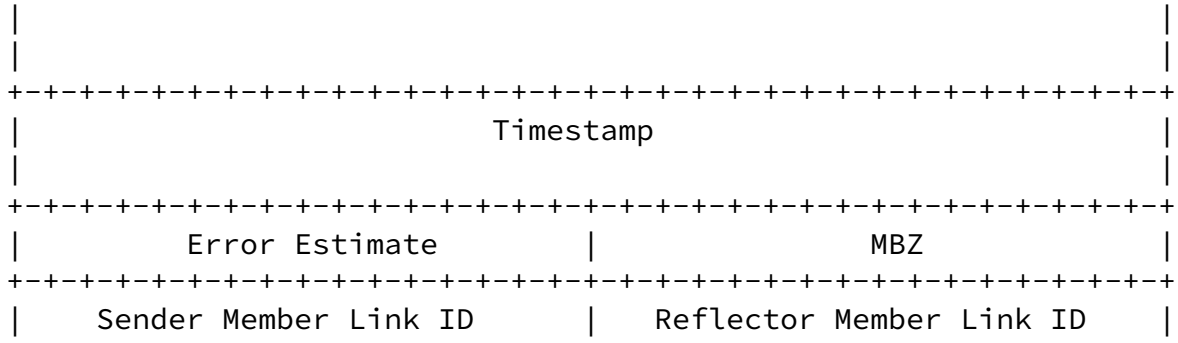
Timestamp
Error Estimate
MBZ
Sender Member Link ID
Reflector Member Link ID

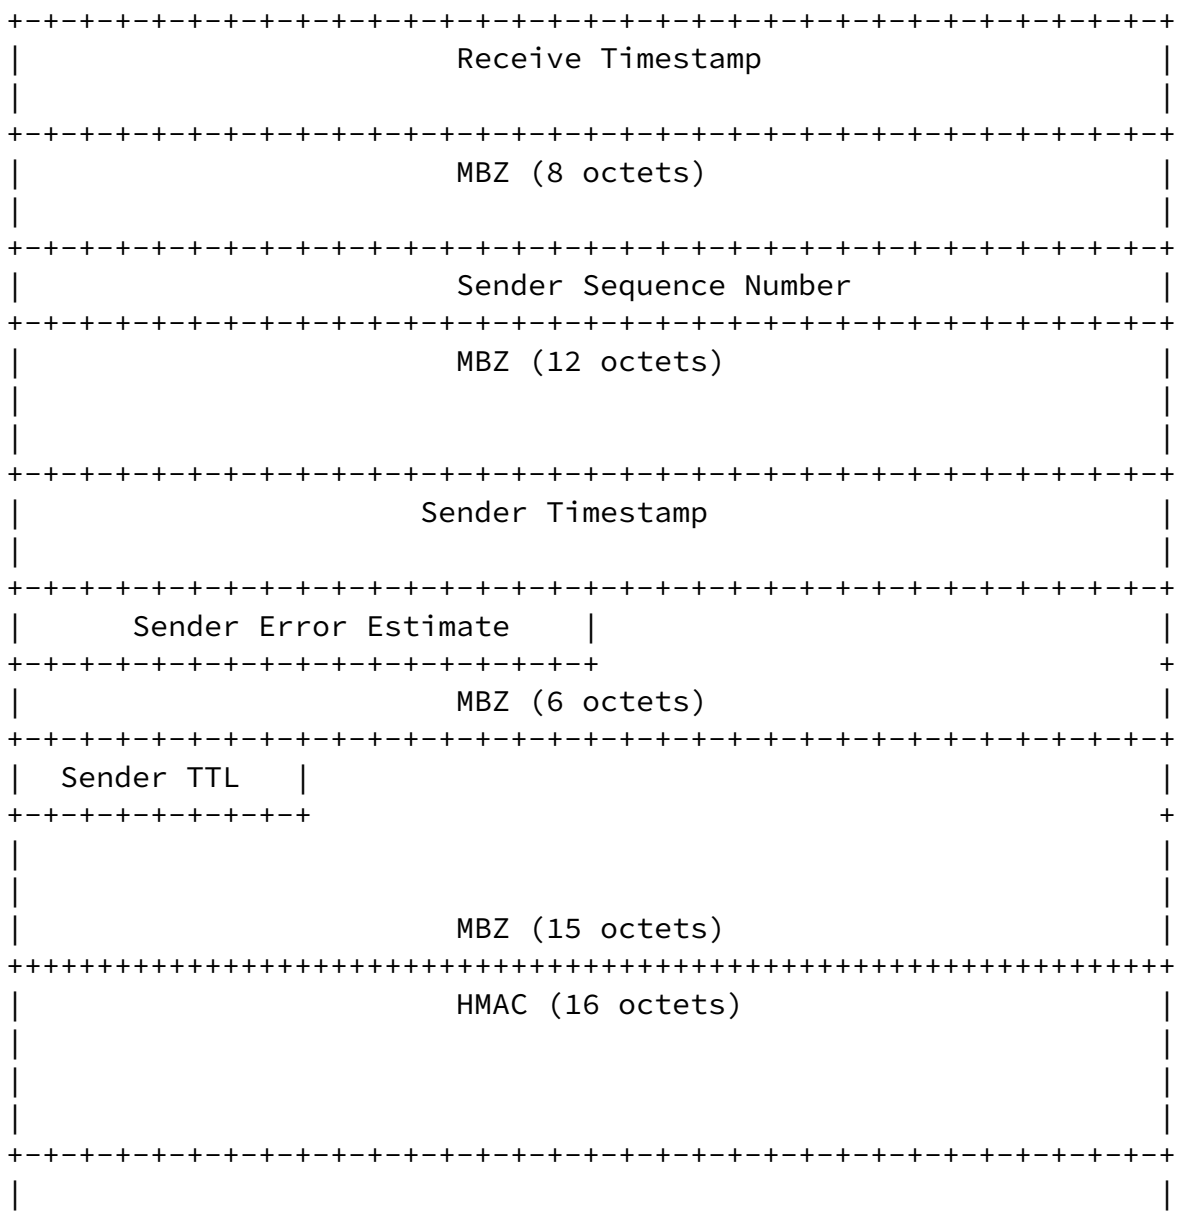
Receive Timestamp
MBZ (8 octets)
Sender Sequence Number
MBZ (12 octets)
Sender Timestamp
Sender Error Estimate
MBZ (6 octets)
Sender TTL
MBZ (15 octets)
HMAC (16 octets)

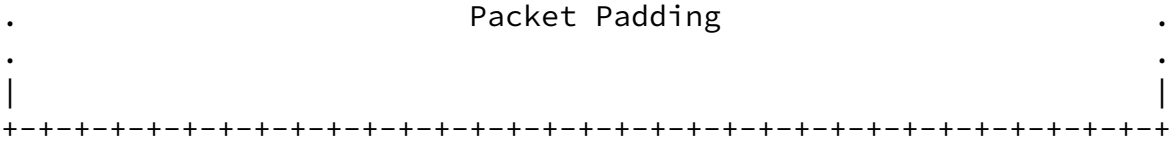

Figure 5: Session-Reflector Packet Format in Authenticated Mode

Except for the Sender/Reflector Member Link ID field, all the other fields are the same as defined in Section 4.2.1 of TWAMP [RFC5357]. Therefore, it follows the same procedure and guidelines as defined in Section 4.2.1 of TWAMP [RFC5357].

Sender Member Link ID (2-octets in length): it is defined to carry the LAG member link identifier of the Sender side. The value of the Sender Member Link ID MUST be unique at the Session-Sender.

Reflector Member Link ID (2-octets in length): it is defined to carry the LAG member link identifier of the Reflector side. The value of the Reflector Member ID MUST be unique at the Session-Reflector.

5. Mirco STAMP Session

5.1. Micro STAMP-Test

The micro STAMP-Test protocol is based on the STAMP-Test protocol [RFC8762] and [I-D.ietf-ippm-stamp-option-tlv] with the following extensions.

5.1.1. Session-Sender Packet Format

The micro STAMP Session-Sender Test packet formats are based on the STAMP Session-Sender Test packet formats and with some extensions, two new fields (Sender and Reflector Member Link ID) are added. The formats are as follows:

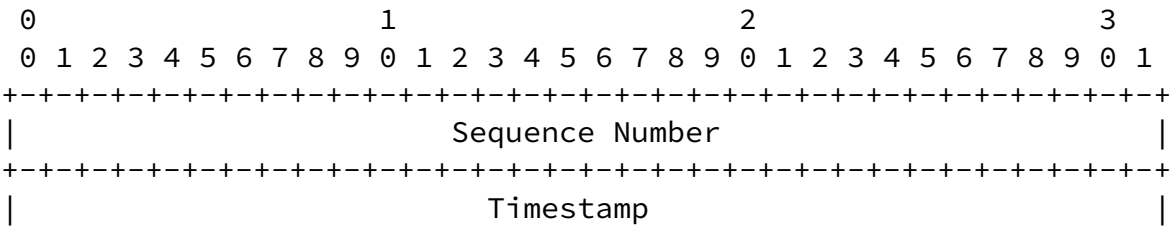

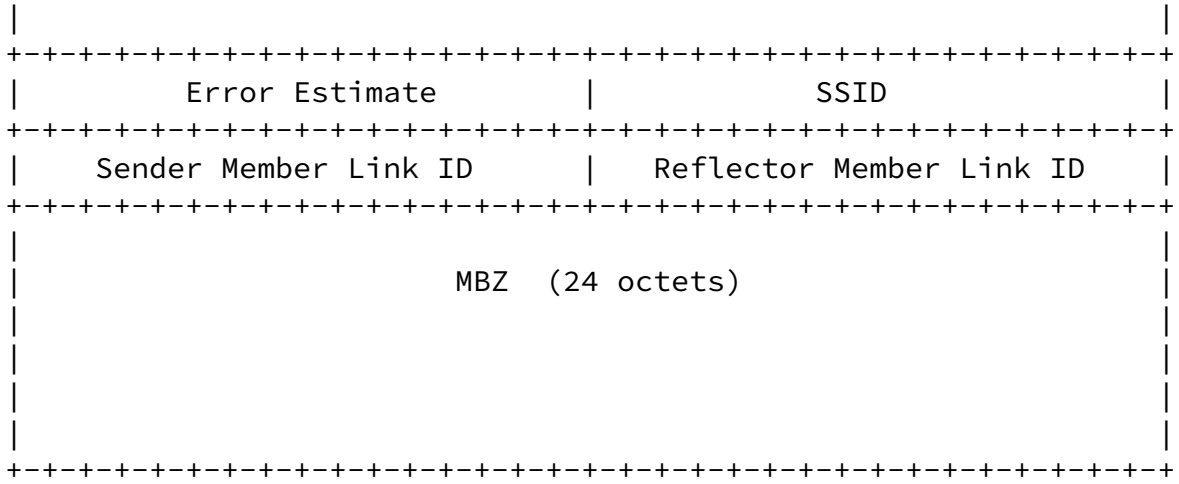

Figure 6: Session-Sender Test Packet in Unauthenticated Mode

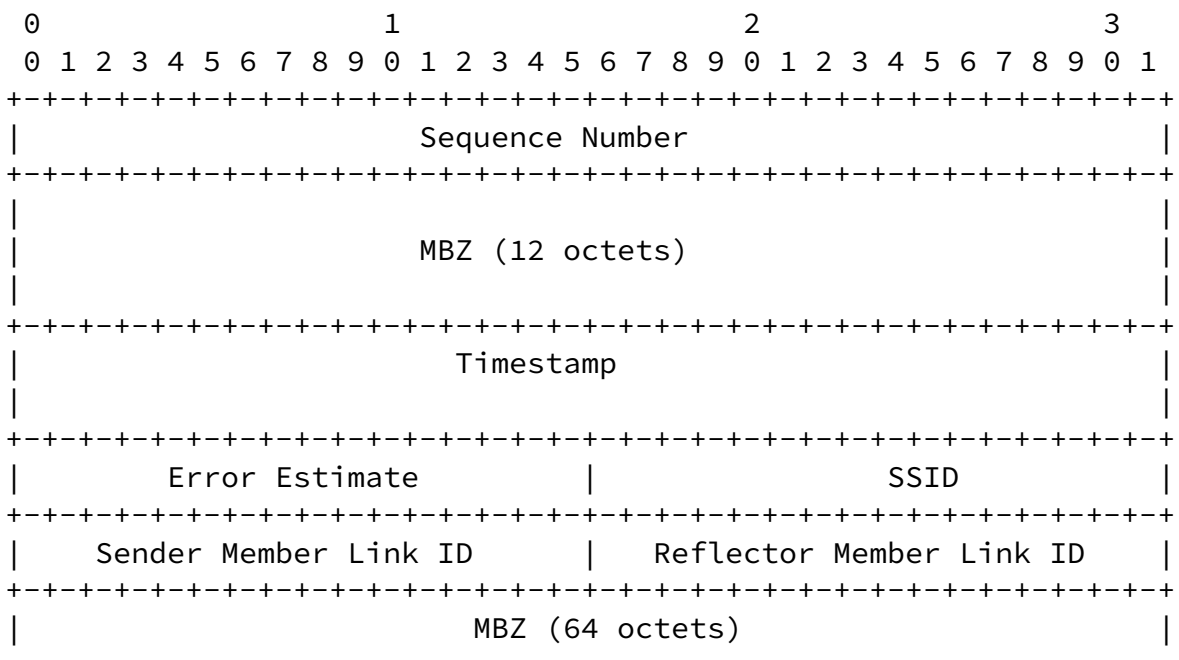

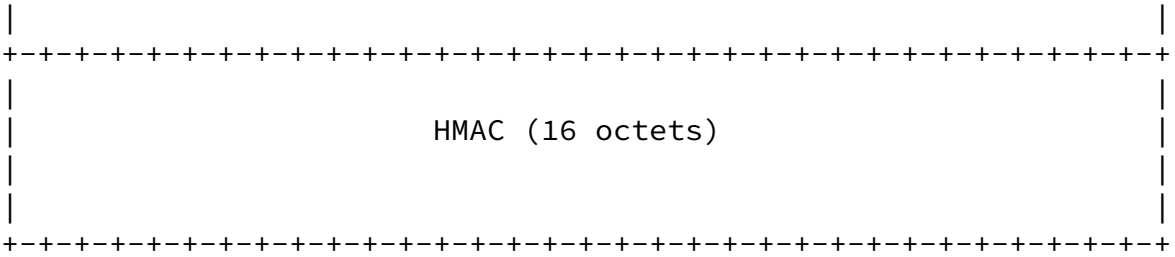

Figure 7: Session-Sender Test Packet in Authenticated Mode

Except for the Sender/Reflector Member Link ID fields, all the other fields are as defined in STAMP [RFC8762] and [I-D.ietf-ippm-stamp-option-tlv].

Sender Member Link ID (2-octets in length): it is defined to carry the LAG member link identifier of the Sender side, which i. The value of the Sender Member Link ID MUST be unique at the Session-Sender.

Reflector Member Link ID (2-octets in length): it is defined to carry the LAG member link identifier of the Reflector side. The value of the Reflector Member ID MUST be unique at the Session-Reflector.

5.1.2. Session-Reflector Packet Format

The micro STAMP Session-Reflector Test packet formats are based on the STAMP Session-Reflector Test packet formats with some minor extensions, two new fields (Sender and Reflector Member Link ID) are added. The formats are as follows:

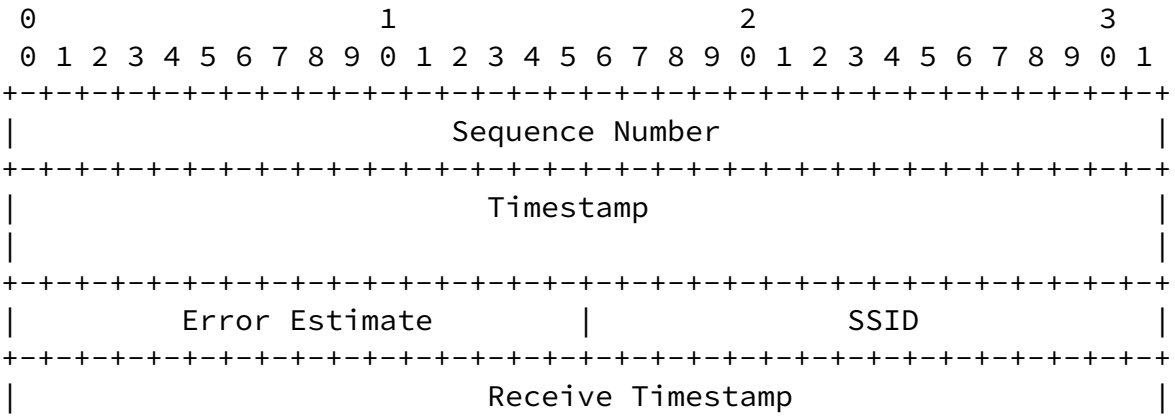

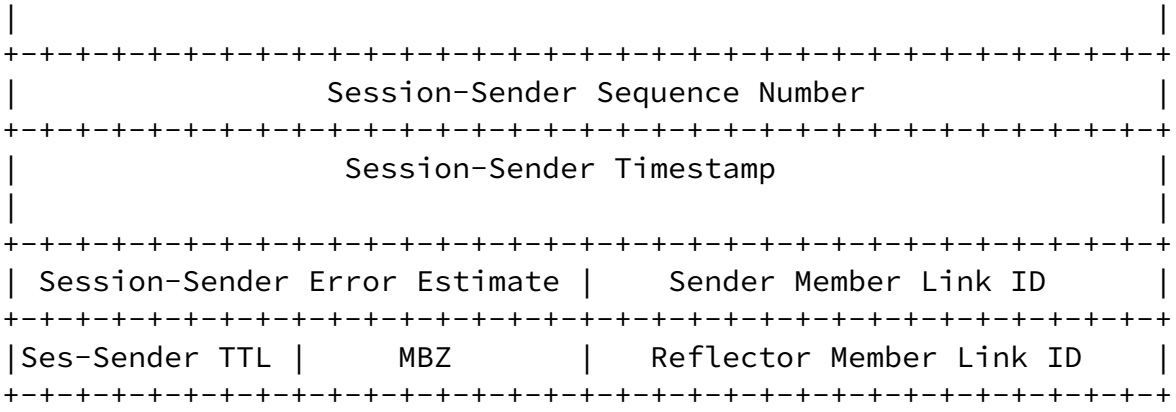

Figure 8: Session-Reflector Test Packet in Unauthenticated Mode

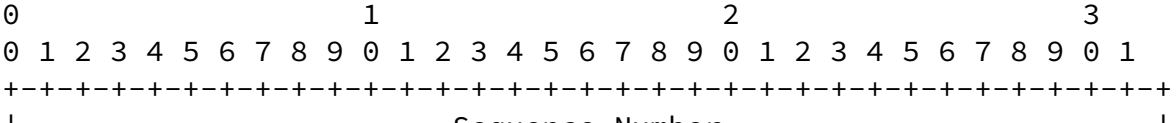

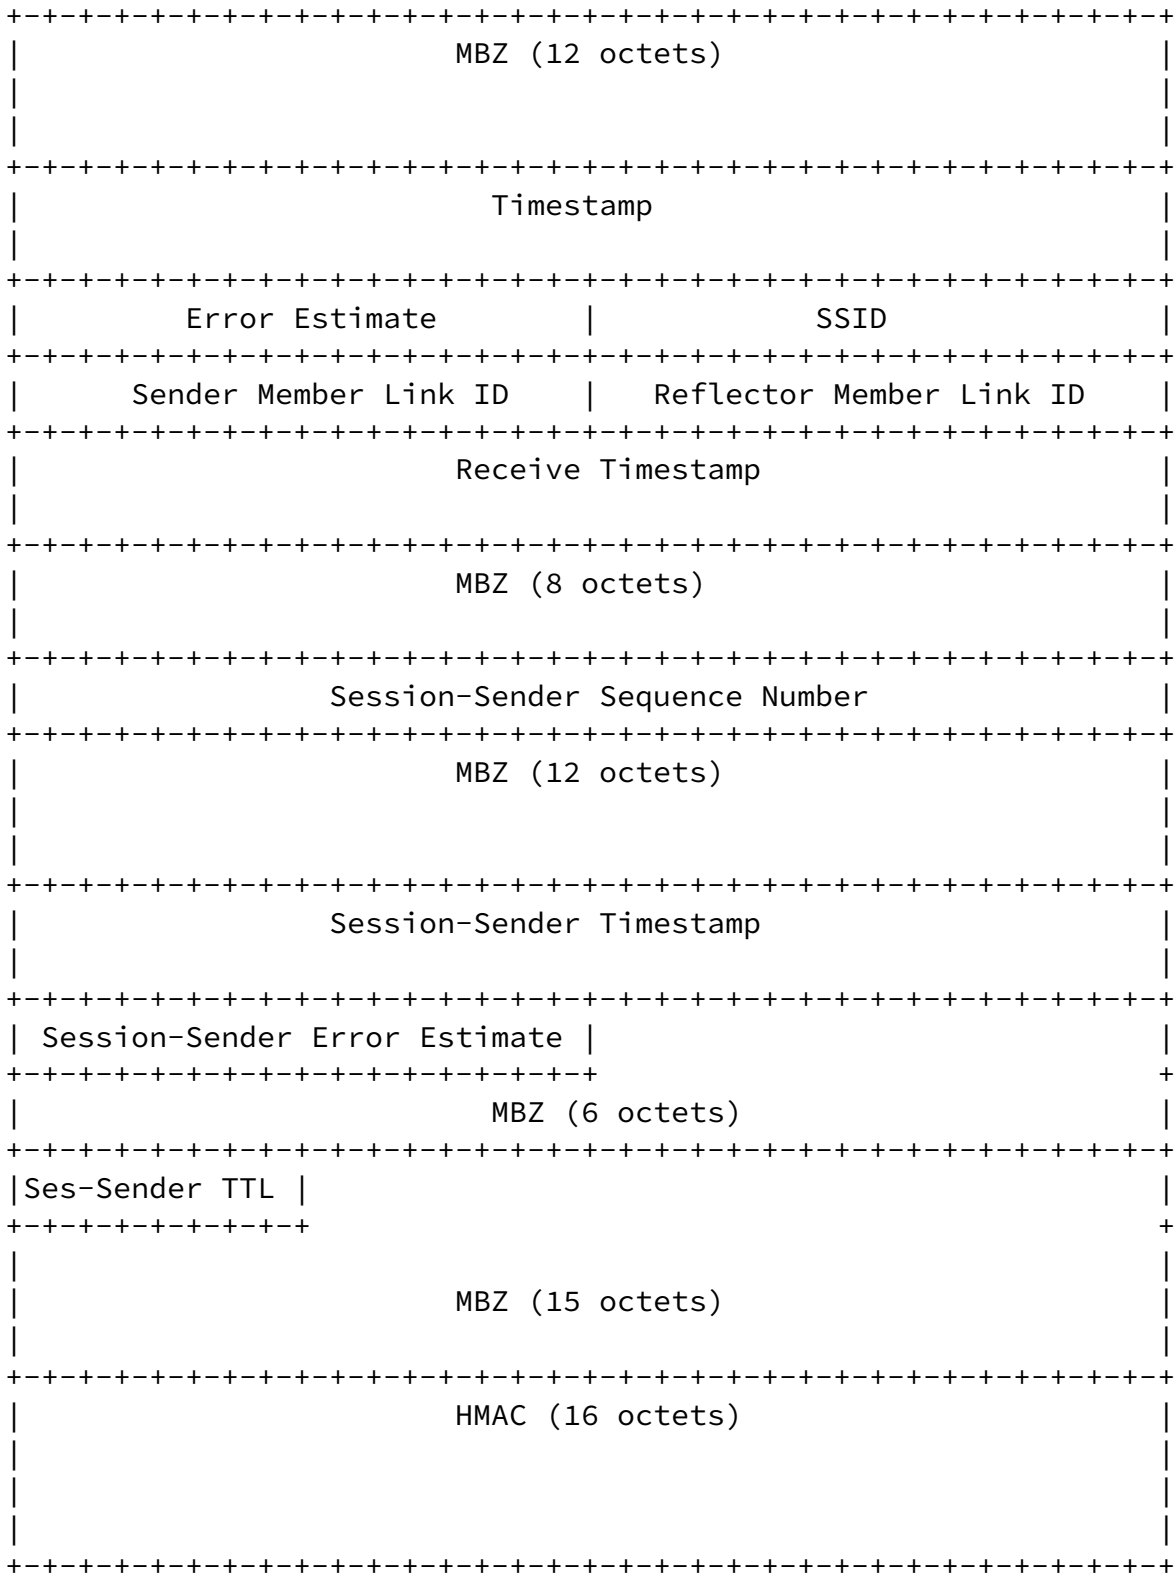

Figure 9: Session-Reflector Test Packet in Authenticated Mode

Except for the Sender/Reflector Member Link ID fields, all the other fields are as defined in STAMP [RFC8762] and [I-D.ietf-ippm-stamp-option-tlv].

Sender Member Link ID (2-octets in length): it is defined to carry the LAG member link identifier of the Sender side. The value of the Sender Member Link ID MUST be unique at the Session-Sender.

Reflector Member Link ID (2-octets in length): it is defined to carry the LAG member link identifier of the Reflector side. The value of the Reflector Member ID MUST be unique at the Session-Reflector.

5.1.3. Micro STAMP-Test Procedures

The micro STAMP-Test reuses the procedures as defined in Section 4 of STAMP [RFC8762] with the following additions:

The micro STAMP Session-Sender MUST send the micro STAMP-Test packets over the member link with which the session is associated.

The configuration and management of the mapping between a micro STAMP session and the Sender/Reflector member link identifiers are outside the scope of this document.

When sending Test packet, the micro STAMP Session-Sender MUST put the Sender member link identifier that is associated with the micro STAMP session in the Sender Member Link ID. If the Session-Sender knows the Reflector member link identifier, it MUST put it in the Reflector Member Link ID fields (see Figure 6 and Figure 7). Otherwise, the Reflector Member Link ID field MUST be set to zero.

The Sender member link identifier is used by the Session-Sender to check whether a reflected Test packet is received from the member link that associates to the correct micro STAMP session. The Reflector member link identifier is used by the Session-Receiver to check whether a Test packet is received from the member link that associates to the correct micro STAMP session.

The Reflector member link identifier can be obtained from pre-configuration or learned through control plane or data plane (e.g., learned from a reflected Test packet). How to abtain/learn the Reflector member link identifier is out of the scope of this document.

When receives a Test packet, the micro STAMP Session-Reflector MUST use the member link from which the Test packet is received to correlate to a micro STAMP session. If there is no such a micro

STAMP session, the Test packet MUST be discarded. If the Reflector

Member Link ID is not zero, the micro STAMP Session-Reflector MUST use the Reflector member link identifier to check whether it associates with the micro STAMP session. If it does not, the Test packet MUST be discarded and no reflected Test packet will be sent back the Session-Sender. If all validation passed, the Session-Reflector sends a reflected Test packet to the Session-Sender. The micro STAMP Session-Reflector MUST put the Sender and Reflector member link identifiers that are associated with the micro STAMP session in the Sender Member Link ID and Reflector Member Link ID fields (see Figure 4 and Figure 9) respectively. The Sender member link identifier is copied from the received Test packet.

When receives a reflected Test packet, the micro STAMP Session-Sender MUST use the member link from which the Test packet is received to correlate to a micro STAMP session. If there is no such a session, the Test packet MUST be discarded. If a matched micro STAMP session exists, the Session-Sender MUST use the Sender Member Link ID to check whether it associates with the session. If the checking failed, the Test packet MUST be discarded.

6. IANA Considerations

6.1. Mico OWAMP-Control Command

This document requires the IANA to allocate the following command type from OWAMP-Control Command Number Registry.

Value	Description	Semantics Definition
TBD1	Request-OW-Micro-Session	This document, Section 3.1

6.2. Mico TWAMP-Control Command

This document requires the IANA to allocate the following command type from TWAMP-Control Command Number Registry.

Value	Description	Semantics Definition
TBD1	Request-TW-Micro-Session	This document, Section 4.1

7. Security Considerations

The security considerations in [RFC4656], [RFC5357], [RFC8762] apply to this document.

8. Acknowledgements

The authors would like to thank Min Xiao, Fang Xin for the valuable comments to this work.

9. References

9.1. Normative References

[I-D.ietf-ippm-stamp-option-tlv]
Mirsky, G., Min, X., Nydell, H., Foote, R., Masputra, A., and E. Ruffini, "Simple Two-way Active Measurement Protocol Optional Extensions", draft-ietf-ippm-stamp-option-tlv-09 (work in progress), August 2020.

[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>.

[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>.

[RFC8762] Mirsky, G., Jun, G., Nydell, H., and R. Foote, "Simple Two-Way Active Measurement Protocol", RFC 8762, DOI 10.17487/RFC8762, March 2020, <https://www.rfc-editor.org/info/rfc8762>.

9.2. Informative References

[I-D.ietf-ippm-ioam-data]
Brockners, F., Bhandari, S., and T. Mizrahi, "Data Fields for In-situ OAM", draft-ietf-ippm-ioam-data-10 (work in progress), July 2020.

[IEEE802.1AX]
IEEE Std. 802.1AX, "IEEE Standard for Local and metropolitan area networks - Link Aggregation", November 2008.

[RFC8321] Fioccola, G., Ed., Capello, A., Cociglio, M., Castaldelli, L., Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi, "Alternate-Marking Method for Passive and Hybrid Performance Monitoring", RFC 8321, DOI 10.17487/RFC8321, January 2018, <https://www.rfc-editor.org/info/rfc8321>.

Authors' Addresses

Zhenqiang Li
China Mobile

Email: li_zhenqiang@hotmail.com

Mach(Guoyi) Chen
Huawei

Email: mach.chen@huawei.com

Greg Mirsky
ZTE Corp.

Email: gregimirsky@gmail.com