

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 23, 2021

Z. Li
China Mobile
M. Chen
Huawei
G. Mirsky
ZTE Corp.
February 19, 2021

**Simple Two-Way Active Measurement Protocol Extensions for Performance
Measurement on LAG
draft-li-ippm-stamp-on-lag-00**

Abstract

This document defines extensions to Simple Two-Way Active Measurement Protocol (STAMP) to implement performance measurement on every member link of a Link Aggregation Group (LAG). Knowing the measured metrics of each member link of a LAG enables operators to enforce a performance metric-based traffic steering policy across 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 August 23, 2021.

Copyright Notice

Copyright (c) 2021 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.	Micro-Session on LAG	3
3.	Micro-STAMP Session	4
3.1.	Micro-STAMP-Test	4
3.1.1.	LAG Member Link ID TLV	4
3.1.2.	Micro-STAMP-Test Procedures	5
4.	IANA Considerations	6
5.	Security Considerations	6
6.	Acknowledgements	6
7.	References	6
7.1.	Normative References	6
7.2.	Informative References	7
	Authors' Addresses	7

[1.](#) Introduction

Link Aggregation Group (LAG), as defined in [[IEEE802.1AX](#)], provides mechanisms to combine multiple physical links into a single logical link. This logical link offers 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.

Usually, when forwarding traffic over a LAG, a hash-based or similar mechanism is used to load-balance the traffic across the LAG member links. In some cases, the link delays of the member links are different because they 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. That requires a solution that could measure the performance metrics of each member link of a LAG.

However, when using Simple Two-Way Active Measurement Protocol (STAMP) [[RFC8762](#)] to measure a LAG's performance, 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. It means that the measured metrics reflect the performance of some member links or an average of some/all member links of the LAG. Therefore, in order to measure every link of a LAG, using active methods would be more appropriate.

This document defines extensions to 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 a LAG (e.g., the LAG includes three member links) directly connects two nodes (A and B). The goal 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 for 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 a 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 the Reflector receives a Test packet, it needs to know from which member link the packet is

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.

3.1.2. 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 associated with the session. The micro STAMP Session-Reflector MUST send the reflected Test packets over the receiving member link.

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

When the Session-Sender sends a Test packet, the LAG Member Link ID TLV MUST be carried, and the Sender member link identifier associated with the micro-STAMP session MUST be put in the Sender Member Link ID field. If the Session-Sender knows the Reflector member link identifier, it MUST set it as the Reflector Member Link ID field's value. Otherwise, the Reflector Member Link ID field MUST be set to zero.

The Session-Sender uses the Sender Member Link ID field's value to check whether the reflected Test packet is received from the member link associated with the correct micro-STAMP session. Therefore, the Session-Reflector MUST copy the Sender Member Link ID value to the reflected Test packet.

The Session-Reflector uses the Reflector Member Link ID value to check whether a Test packet is received from the member link associated with the correct micro-STAMP session.

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

When the micro-STAMP Session-Reflector receives a Test packet, it MUST use the receiving member link to correlate the Test packet to a micro-STAMP session. If there is no such a micro-STAMP session, the Test packet MUST be discarded. Suppose the Reflector Member Link ID is not zero. In that case, the micro-STAMP Session-Reflector MUST use the Reflector member link identifier to check whether it is associated with the micro-STAMP session. If it is not, the Test packet MUST be discarded and no reflected Test packet will be sent back to the Session-Sender. If all validation passed, the Session-Reflector sends a reflected Test packet to the Session-Sender over the receiving member link. The micro-STAMP Session-Reflector MUST

put the Sender and Reflector member link identifiers associated with the micro-STAMP session in the Sender Member Link ID and Reflector Member Link ID fields, respectively. The Sender member link identifier is copied from the received Test packet.

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

4. IANA Considerations

This document requires the IANA to allocate the following the TLV type from the "STAMP TLV Types" sub-registry.

Value	Description	Reference
TBD1	LAG Member Link ID	This document

5. Security Considerations

This document does not introduce additional security requirements and mechanisms other than the ones described in [RFC8762] apply to this document.

6. Acknowledgements

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

7. References

7.1. Normative References

- [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>>.
- [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>>.
- [RFC8972] Mirsky, G., Min, X., Nydell, H., Foote, R., Masputra, A., and E. Ruffini, "Simple Two-Way Active Measurement Protocol Optional Extensions", [RFC 8972](#), DOI 10.17487/RFC8972, January 2021, <<https://www.rfc-editor.org/info/rfc8972>>.

7.2. 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-11](#) (work in progress), November 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

