Internet Engineering Task Force (IETF)
Request for Comments: 9534
Category: Standards Track
ISSN: 2070-1721

Z. Li
China Mobile
T. Zhou
Huawei
J. Guo
ZTE Corp.
G. Mirsky
Ericsson
R. Gandhi
Cisco Systems, Inc.

January 2024

Simple Two-Way Active Measurement Protocol Extensions for Performance Measurement on a Link Aggregation Group

Abstract

This document extends 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-based traffic steering policy across the member links.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at https://www.rfc-editor.org/info/rfc9534.

Table of Contents

1. Introduction

A 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 mechanism is used to load balance the traffic across the LAG member links. The link delay might vary between member links because of different transport paths, especially when a LAG is used in a wide area network. To provide low-latency service for time-sensitive traffic, we need to explicitly steer the traffic across the LAG member links based on the link delay, loss, and so on. That requires a solution to measure the performance metrics of each member link of a LAG. Hence, the measured performance metrics can work together with Layer 2 bundle member link attributes advertisement [RFC8668] for traffic steering.

According to the classifications in [RFC7799], Simple Two-way Active Measurement Protocol (STAMP) [RFC8762] is an active measurement method, and it can complement passive and hybrid methods. It provides a mechanism to measure both one-way and round-trip performance metrics, like delay, delay variation, and packet loss. A STAMP test session over the LAG can be used to measure the performance of a member link using a specially constructed 5-tuple. The session can be used to measure an average of some or all member links of the LAG by varying one or more elements of that 5-tuple. However, without the knowledge of each member link, a STAMP test session cannot measure the performance of every physical member link.

This document extends STAMP to implement performance measurement on every member link of a LAG. It can provide the same metrics as One-Way Active Measurement Protocol (OWAMP) [RFC4656] and Two-Way Active Measurement Protocol (TWAMP) [RFC5357] can measure, such as delay, jitter, and packet loss.

1.1. 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.

2. Micro Sessions on a LAG

This document addresses the scenario where a LAG directly connects two nodes. An example of this is in Figure 1, where the LAG consisting of four links connects nodes A and B. The goal is to measure the performance of each link of the LAG.

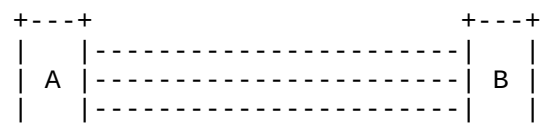

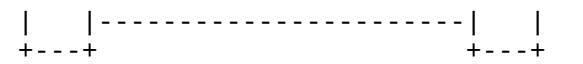

Figure 1: Performance Measurement on a LAG

To measure the performance metrics of every member link of a LAG, multiple sessions (one session for each member link) need to be established between the two endpoints that are connected by the LAG. These sessions are called "micro sessions" in the remainder of this document. Although micro sessions are in fact STAMP sessions established on member links of a LAG, test packets of micro sessions MUST carry member link information for validation.

All micro sessions of a LAG share the same Sender IP Address and Receiver IP Address. As for the UDP port, the micro sessions may share the same Sender Port and Receiver Port pair or each micro session may be configured with a different Sender Port and Receiver Port pair. From the operational point of view, the former is simpler and is RECOMMENDED.

Test packets of a micro session MUST carry the member link information for validation checks. For example, when a micro STAMP Session-Sender receives a reflected test packet, it checks whether the test packet is from the expected member link. The member link information is encoded in the Micro-session ID TLV introduced in Section 3, which also provides a detailed description about member link validation.

A micro STAMP Session-Sender MAY include the Follow-Up Telemetry TLV [RFC8972] to request information from the micro Session-Reflector. This timestamp might be important for the micro Session-Sender, as it improves the accuracy of network delay measurement by minimizing the impact of egress queuing delays on the measurement.

3. Member Link Validation

Test packets MUST carry member link information in the Micro-session ID TLV introduced in this section for validation checks. The micro Session-Sender verifies whether the test packet is received from the expected member link. It also verifies whether the packet is sent from the expected member link at the Reflector side. The micro Session-Reflector verifies whether the test packet is received from the expected member link.

3.1. Micro-session ID TLV

The STAMP TLV mechanism [RFC8972] extends STAMP test packets with one or more optional TLVs. This document defines the TLV Type (value 11) for the Micro-session ID TLV that carries the micro STAMP Session-Sender member link identifier and Session-Reflector member link identifier in the Sender Micro-session ID field and the Reflector Micro-session ID field, respectively. The format of the Micro-session ID TLV is shown as follows:

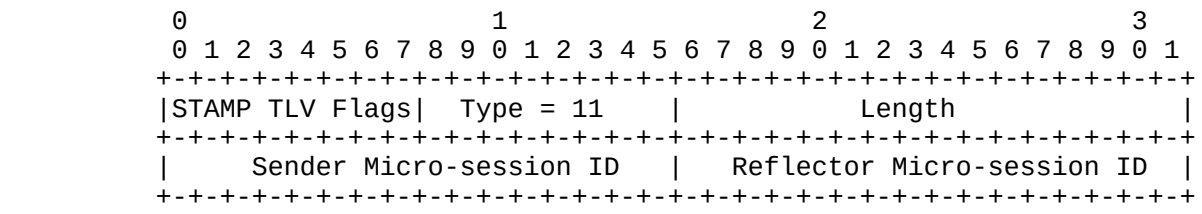

Figure 2: Micro-session ID TLV

Type (1 octet in length): This field is defined to indicate this TLV is a Micro-session ID TLV. Value 11 has been allocated by IANA

(Section 5).

Length (2 octets in length): This field is defined to carry the length of the Value field in octets. The Length field value MUST be 4.

Sender Micro-session ID (2 octets in length): This field is defined to carry the LAG member link identifier of the Sender side. In the future, it may be used generically to cover use cases beyond LAGs. The value of this field MUST be unique within a STAMP session at the Session-Sender.

Reflector Micro-session ID (2 octets in length): This field is defined to carry the LAG member link identifier of the Reflector side. In the future, it may be used generically to cover use cases beyond LAGs. The value of this field MUST be unique within a STAMP session at the Session-Reflector.

3.2. 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 with which the session is associated. The mapping between a micro STAMP session and the Sender/Reflector member link identifiers can be configured by augmenting the STAMP YANG [STAMP-YANG]. The detailed augmentation is not in the scope of this document.

When sending a test packet, the micro STAMP Session-Sender MUST set the Sender Micro-session ID field with the member link identifier associated with the micro STAMP session. If the Session-Sender knows the Reflector member link identifier, the Reflector Micro-session ID field MUST be set. Otherwise, the Reflector Micro-session ID field MUST be zero. The Reflector member link identifier can be obtained from preconfiguration or learned from data plane (e.g., the reflected test packet). This document does not specify the way to obtain the Reflector member link identifier.

When the micro STAMP Session-Reflector receives a test packet, if the Reflector Micro-session ID is not zero, the micro STAMP Session-Reflector MUST use the Reflector member link identifier to check whether it is associated with the micro STAMP session. If the validation fails, the test packet MUST be discarded. If the Reflector Micro-session ID is zero, it will not be verified. If all validations 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 Micro-session ID and Reflector Micro-session ID fields, respectively. The Sender member link identifier is copied from the received test packet.

When receiving a reflected test packet, the micro Session-Sender MUST use the Sender Micro-session ID to validate whether the reflected test packet is correctly received from the expected member link. If the validation fails, the test packet MUST be discarded. The micro Session-Sender MUST use the Reflector Micro-session ID to validate the Reflector’s behavior. If the validation fails, the test packet MUST be discarded.

Two modes of the STAMP Session-Reflector, stateless and stateful, characterize the expected behavior as described in Section 4 of STAMP [RFC8762]. The micro STAMP-Test also supports both stateless and stateful modes. However, the micro STAMP-Test does not introduce any additional state to STAMP, i.e., any procedure with regard to the

Micro-session ID is stateless.

4. Applicability

The micro STAMP Session-Sender sends micro Session-Sender packets with the Micro-session ID TLV. The micro Session-Reflector checks whether a test packet is received from the member link associated with the correct micro STAMP session if the Reflector Micro-session ID field is set. When reflecting, the micro STAMP Session-Reflector copies the Sender Micro-session ID from the received micro Session-Sender packet to the micro Session-Reflector packet and sets the Reflector Micro-session ID field with the member link identifier that is associated with the micro STAMP session. When receiving the micro Session-Reflector packet, the micro Session-Sender uses the Sender Micro-session ID to check whether the packet is received from the member link associated with the correct micro STAMP session. The micro Session-Sender also use the Reflector Micro-session ID to validate the Reflector’s behavior.

5. IANA Considerations

IANA has allocated the following STAMP TLV Type for the Micro-session ID TLV in the "STAMP TLV Types" registry [RFC8972]:

Value	Description	Reference
11	Micro-session ID	This Document

Table 1: New STAMP TLV Type

6. Security Considerations

The STAMP extension defined in this document is intended for deployment in the LAG scenario where Session-Sender and Session-Reflector are directly connected. As such, it’s assumed that a node involved in a STAMP operation has previously verified the integrity of the LAG connection and the identity of its one-hop-away peer node.

This document does not introduce any additional security issues, and the security mechanisms defined in [RFC8762] and [RFC8972] apply in this document.

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

[IEEE802.1AX]
IEEE, "IEEE Standard for Local and Metropolitan Area Networks -- Link Aggregation", IEEE Std 802.1AX-2020, DOI 10.1109/IEEESTD.2020.9105034, May 2020, <https://ieeexplore.ieee.org/document/9105034>.

[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>.

[RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <https://www.rfc-editor.org/info/rfc7799>.

[RFC8668] Ginsberg, L., Ed., Bashandy, A., Filsfils, C., Nanduri, M., and E. Aries, "Advertising Layer 2 Bundle Member Link Attributes in IS-IS", RFC 8668, DOI 10.17487/RFC8668, December 2019, <https://www.rfc-editor.org/info/rfc8668>.

[STAMP-YANG]
Mirsky, G., Min, X., Luo, W. S., and R. Gandhi, "Simple Two-way Active Measurement Protocol (STAMP) Data Model", Work in Progress, Internet-Draft, draft-ietf-ippm-stamp-yang-12, 5 November 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-ippm-stamp-yang-12>.

Acknowledgements

The authors would like to thank Mach Chen, Min Xiao, Fang Xin, Marcus Ihlar, and Richard Foote for the valuable comments to this work.

Authors' Addresses

Zhenqiang Li
China Mobile
No. 29 Finance Avenue
Xicheng District
Beijing
China
Email: li_zhenqiang@hotmail.com

Tianran Zhou
Huawei
China
Email: zhoutianran@huawei.com

Jun Guo
ZTE Corp.
China
Email: guo.jun2@zte.com.cn

Greg Mirsky

Ericsson
United States of America
Email: gregimirsky@gmail.com

Rakesh Gandhi
Cisco Systems, Inc.
Canada
Email: rgandhi@cisco.com