Internet Engineering Task Force (IETF)
Request for Comments: 9533
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

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

Abstract

This document defines extensions to the One-Way Active Measurement Protocol (OWAMP) and the Two-Way Active Measurement Protocol (TWAMP) 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 the 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/rfc9533.

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 every 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], OWAMP [RFC4656] and TWAMP [RFC5357] are active measurement methods, and they can complement passive and hybrid methods. With either method, one 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 test session cannot measure the performance of every physical member link.

This document extends OWAMP and TWAMP to implement performance measurement on every member link of a LAG. It can provide the same metrics as OWAMP and TWAMP 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.

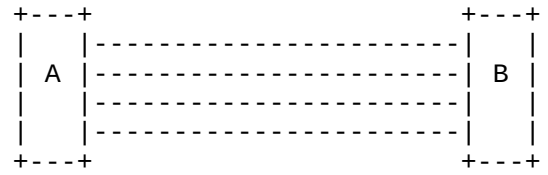

Figure 1: Performance Measurement on a LAG

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 OWAMP or TWAMP sessions established on member links of a LAG, test packets of micro TWAMP 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 TWAMP Session-Sender receives a reflected test packet, it checks whether the test packet is from the expected member link.

3. Micro OWAMP Session

3.1. Micro OWAMP-Control

To support the micro OWAMP session, a new command, Request-OW-Micro-Sessions (5), is defined in this document. The Request-OW-Micro-Sessions command is based on the OWAMP Request-Session command and uses the message format as described in Section 3.5 of [RFC4656]. Test session creation of micro OWAMP sessions follows the same procedure as defined in Section 3.5 of [RFC4656] with the following additions:

When an OWAMP Server receives a Request-OW-Micro-Sessions command, if the request is accepted, the OWAMP Server MUST build a set of micro sessions for all the member links of the LAG from which the Request-OW-Micro-Sessions 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 [RFC4656] with the following additions:

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

4. Micro TWAMP Session

4.1. Micro TWAMP-Control

To support the micro TWAMP session, a new command, Request-TW-Micro-

Sessions (11), is defined in this document. The Request-TW-Micro-Sessions command is based on the TWAMP Request-Session command and uses the message format as described in Section 3.5 of [RFC5357]. Test session creation of micro TWAMP sessions follows the same procedure as defined in Section 3.5 of [RFC5357] with the following additions:

When a TWAMP Server receives a Request-TW-Micro-Sessions command, if the request is accepted, the TWAMP Server MUST build a set of micro sessions for all the member links of the LAG from which the Request-TW-Micro-Sessions message is received.

4.2. Micro TWAMP-Test

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

4.2.1. Sender 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]. Two new fields (Sender Micro-session ID and Reflector Micro-session ID) are added to carry the LAG member link identifiers.

For unauthenticated mode, the format is as below:

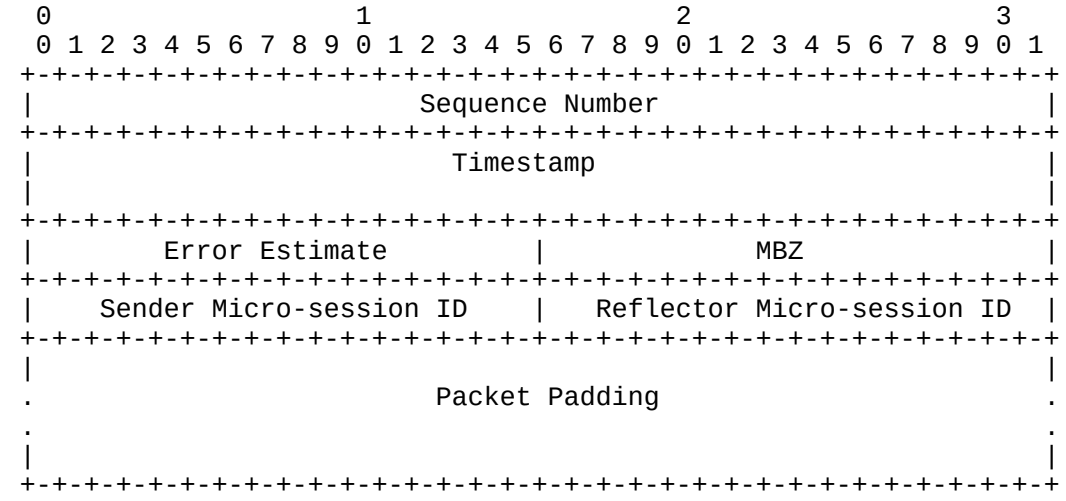

Figure 2: Micro Session-Sender Packet Format in Unauthenticated Mode

For authenticated and encrypted mode, the format is as below:

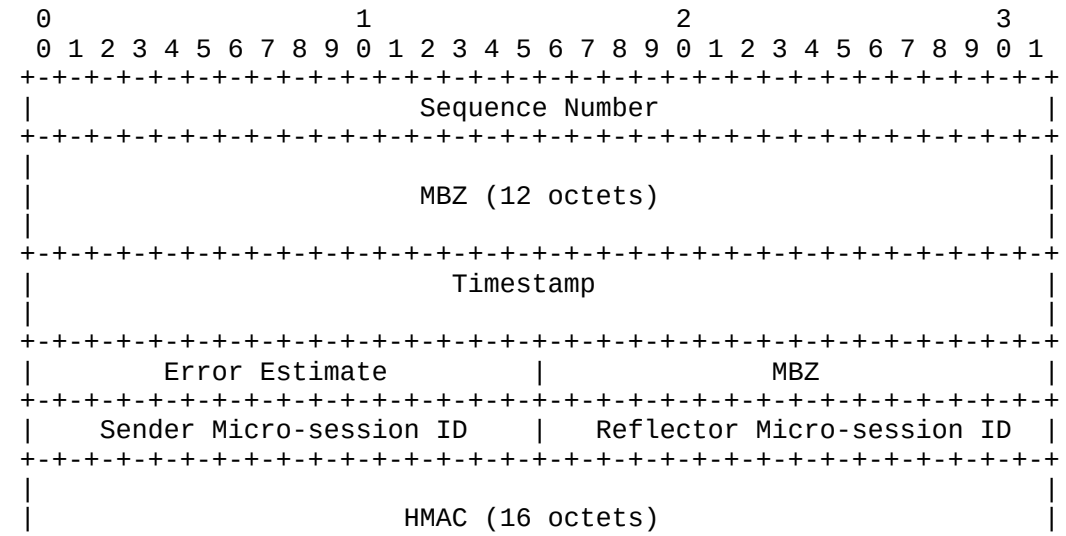

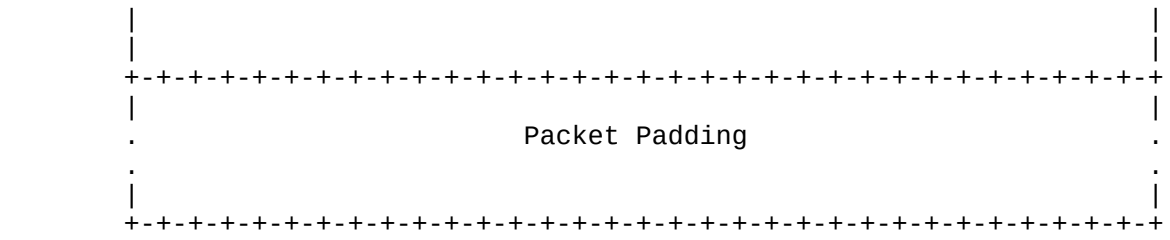

Figure 3: Micro Session-Sender Packet Format in Authenticated Mode

Except for the Sender Micro-session ID field and the Reflector Micro-session ID field, all the other fields are the same as defined in Section 4.1.2 of [RFC5357] and follow the procedure and guidelines defined therein.

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 TWAMP 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 TWAMP session at the Session-Reflector.

4.2.2. Sender Behavior

The micro TWAMP Session-Sender inherits the behaviors of the TWAMP Session-Sender as defined in Section 4.1 of [RFC5357]. In addition, the micro TWAMP Session-Sender MUST send the micro Session-Sender test packets over the member link with which the session is associated.

When sending the 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 Micro-session ID. If the Session-Sender knows the Reflector member link identifier, the Reflector Micro-session ID field (see Figures 2 and 3) MUST be set. Otherwise, the Reflector Micro-session ID field MUST be zero.

A test packet with a Sender member link identifier is sent to the Session-Reflector and then is reflected with the same Sender member link identifier. So the Session-Sender can use the Sender member link identifier to check whether a reflected test packet is received from the member link associated with the correct micro TWAMP session.

The Reflector member link identifier carried in the Reflector Micro-session ID field is used by the Session-Reflector to check whether a test packet is received from the member link associated with the correct micro TWAMP session. It means that the Session-Sender has to learn the Reflector member link identifier. Once the Session-Sender knows the Reflector member link identifier, it MUST put the identifier in the Reflector Micro-session ID field (see Figures 2 or 3) of the test packets that will be sent to the Session-Reflector. The Reflector member link identifier can be obtained from preconfiguration or learned from the data plane (e.g., the reflected test packet). This document does not specify the way to obtain the Reflector member link identifier.

When receiving a reflected test packet, the micro TWAMP Session-Sender MUST use the receiving member link to correlate the reflected test packet to a micro TWAMP session. If there is no such session, the reflected test packet MUST be discarded. If a matched session

exists, 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.

4.2.3. Reflector 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]. Two new fields (Sender and Reflector Micro-session ID) are added to carry the LAG member link identifiers.

For unauthenticated mode, the format is as below:

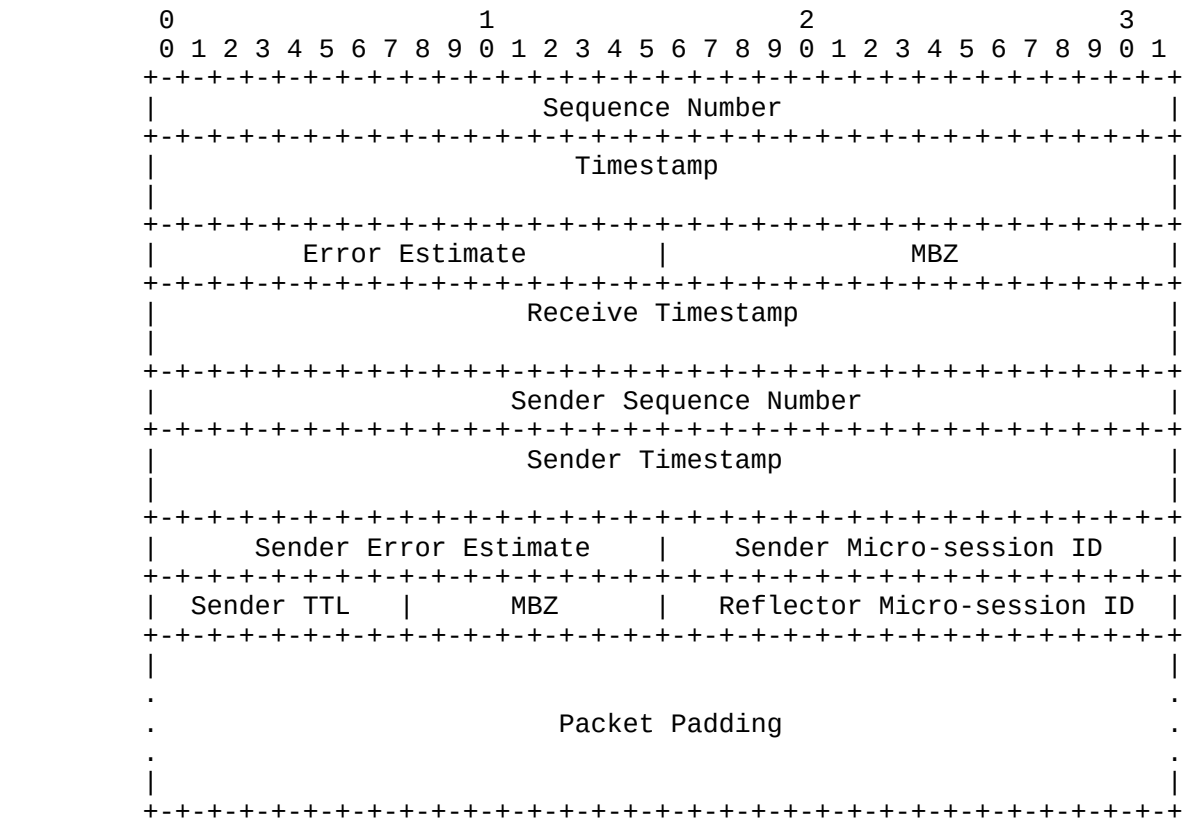

Figure 4: Micro Session-Reflector Packet Format in Unauthenticated Mode

For authenticated and encrypted mode, the format is as below:

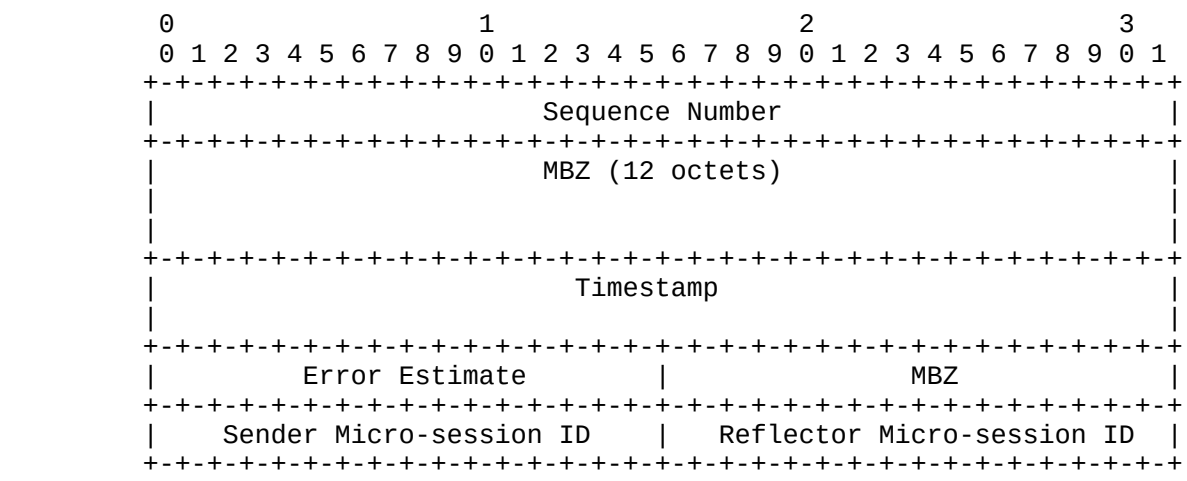

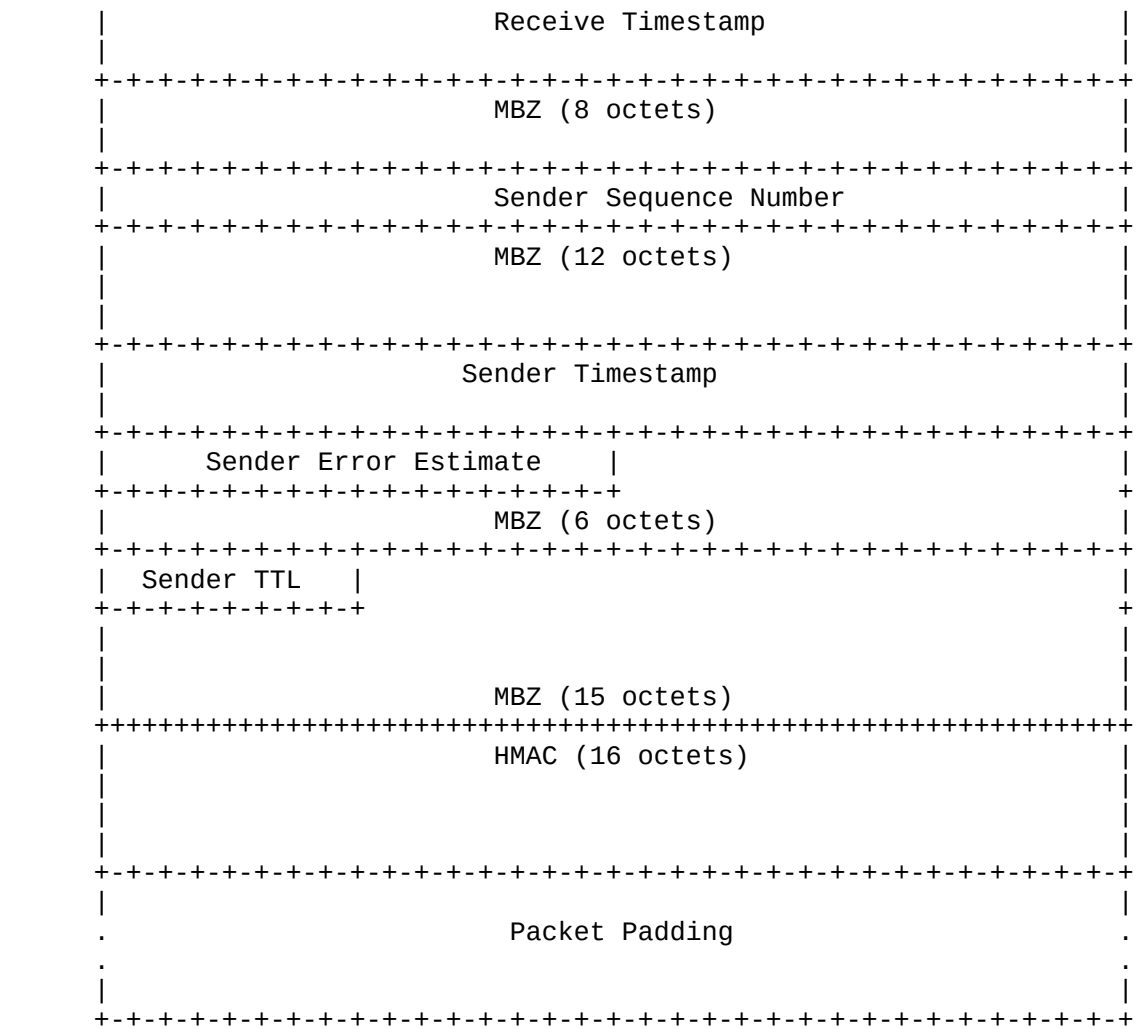

Figure 5: Micro Session-Reflector Packet Format in Authenticated Mode

Except for the Sender Micro-session ID field and the Reflector Micro-session ID field, all the other fields are the same as defined in Section 4.2.1 of [RFC5357] and follow the same procedure and guidelines defined therein.

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 TWAMP 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 TWAMP session at the Session-Reflector.

4.2.4. 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 receiving a test packet, the micro TWAMP Session-Reflector MUST use the receiving member link to correlate the test packet to a micro TWAMP session. If there is no such a session, the test packet MUST be discarded. If the Reflector Micro-session ID is not zero, the Reflector MUST use the Reflector Micro-session ID to validate whether it associates with the receiving member link. If

the Reflector Micro-session ID is zero, it will not be verified. If the validation fails, the test packet MUST be discarded.

When sending a response to the received test packet, the micro TWAMP Session-Reflector MUST copy the Sender member link identifier from the received test packet and put it in the Sender Micro-session ID field of the reflected test packet (see Figures 4 and 5). In addition, the micro TWAMP Session-Reflector MUST fill the Reflector Micro-session ID field (see Figures 4 and 5) of the reflected test packet with the member link identifier that is associated with the micro TWAMP session.

5. Applicability

To set up the micro OWAMP sessions, the Control-Client sends the Request-OW-Micro-Sessions command to the OWAMP Server. The OWAMP Server accepts the request and builds a set of micro sessions for all the member links of the LAG.

For micro TWAMP sessions, a similar set up procedure is used. Then, the micro TWAMP Session-Sender sends micro Session-Sender packets with the Sender Micro-session ID and the Reflector Micro-session ID. If the Reflector Micro-session ID field is set, the micro Session-Reflector checks whether a test packet is received from the member link associated with the correct micro TWAMP session. When reflecting, the micro TWAMP Session-Reflector copies the Sender Micro-session ID from the received micro Session-Sender packet to the micro Session-Reflector packet; then, it sets the Reflector Micro-session ID field with the member link identifier that is associated with the micro TWAMP session. When receiving the micro TWAMP 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 TWAMP session. The micro Session-Sender also uses the Reflector Micro-session ID to validate the Reflector’s behavior.

6. IANA Considerations

6.1. Micro OWAMP-Control Command

IANA has allocated the following command type from the "OWAMP-Control Command Numbers" registry.

Value	Description	Reference
5	Request-OW-Micro-Sessions	This document

Table 1: Request-OW-Micro-Sessions Command Number

6.2. Micro TWAMP-Control Command

IANA has allocated the following command type from the "TWAMP-Control Command Numbers" registry.

Value	Description	Reference
11	Request-TW-Micro-Sessions	This document

Table 2: Request-TW-Micro-Sessions Command Number

7. Security Considerations

This document does not introduce additional security requirements and mechanisms other than those described in [RFC4656] and [RFC5357].

8. References

8.1. 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>.

[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>.

[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>.

8.2. 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>.

[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>.

Acknowledgements

The authors would like to thank Fang Xin, Henrik Nydell, Mach Chen, Min Xiao, Jeff Tantsura, 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