IPPM Working Group
Internet-Draft
Intended status: Standards Track
Expires: 7 August 2024

R. Gandhi, Ed.
C. Filsfils
Cisco Systems, Inc.
D. Voyer
Bell Canada
M. Chen
Huawei
B. Janssens
Colt

4 February 2024

Simple Two-Way Direct Loss Measurement Procedure
draft-gandhi-ippm-simple-direct-loss-07

Abstract

This document defines Simple Two-Way Direct Loss Measurement (DLM) procedure that can be used for Alternate-Marking Method for detecting accurate data packet loss in a network. Specifically, DLM probe packets are defined for both unauthenticated and authenticated modes and they are efficient for hardware-based implementation.

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 7 August 2024.

Copyright Notice

Table of Contents

1. Introduction

Many Service Provider Service Level Agreements (SLAs) depend on the ability to measure performance loss metric experienced by the Customer data traffic flow. Accurate Customer data packet loss can be measured by using a Direct Loss Measurement (DLM) procedure. Currently there is no efficient active measurement procedure available for accurate data packet loss detection in IP networks. Note that an approach for conducting packet loss measurement in an IP network is documented in [RFC7680]. This approach requires clock synchronization between the measurement points and lacks support for accurate data packet loss measurement.

[ITU-Y1731] defines procedures for performance loss monitoring for Ethernet-based networks. Specifically, the Loss Measurement Message (LMM) defined in Section 9.12 of [ITU-Y1731] can be used for accurate frame loss measurement as described in Appendix II of that document. The procedure is specific to the Ethernet-based networks and does not apply to the IP networks.

The Simple Two-Way Active Measurement Protocol (STAMP) provides capabilities for the measurement of various performance metrics in IP networks [RFC8762] without the use of a control channel to pre-signal session parameters. The STAMP can be used for (synthetic or inferred) packet loss measurement based on the Sequence Number in the test packets, however, this method can only provide approximate packet loss metrics.

[RFC8972] defines optional extensions for STAMP. The STAMP test packet with the "Direct Measurement" TLV (Type 5) [RFC8972] can be used for combined timestamps and data packet counters collection. This method, however, has the following limitations when used for detecting data packet loss:

* For hardware-based implementation (e.g., in an ASIC), the optional "Direct Measurement" TLV adds unnecessary processing overhead on the Session-Reflector as not all STAMP Session-Sender test packets carry the "Direct Measurement" TLV and also there can be multiple TLV Types present. This also means that the location of the transmit counter is not at the fixed location in the STAMP test packets.

* The STAMP "Direct Measurement" TLV does not support 64-bit counters, counters for bytes, counters per traffic class.

* The STAMP "Direct Measurement" TLV also does not identify the Block Number of the Direct Measurement, which is required for Alternate-Marking Method (AMM) [RFC9341] for data packet loss measurement.

This document defines Simple Two-Way Direct Loss Measurement (DLM) procedure that can be used for Alternate-Marking Method [RFC9341] for detecting accurate data packet loss in a network. Specifically, DLM probe packets are defined for both unauthenticated and authenticated modes and they are efficient for hardware-based implementation (e.g., in an ASIC).

2. Conventions Used in This Document

2.1. 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.

2.2. Abbreviations

AMM: Alternate-Marking Method.

DLM: Direct Loss Measurement.

HMAC: Hashed Message Authentication Code.

MBZ: Must be Zero.

PM: Performance Measurement.

SHA: Secure Hash Algorithm.

SSID: Sender Session Identifier.

STAMP: Simple Two-Way Active Measurement Protocol.

TTL: Time To Live.

2.3. Reference Topology

As shown in the Reference Topology, the Session-Sender S1 initiates a Direct Loss Measurement (DLM) probe packet over IP/UDP transport. The Session-Reflector R1 receives the Session-Sender's DLM probe packet and acts according to the local configuration. The Session-Reflector R1 transmits a DLM reply probe packet to the Session-Sender S1. The C1 is a transmit counter and C4 is a receive counter added by node S1. The C2 is a receive counter and C3 is a transmit counter added by node R1.

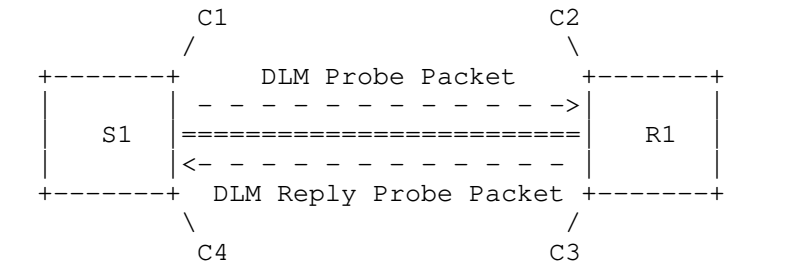

Reference Topology

3. Overview

For accurate data packet loss detection, the DLM probe packets are transmitted by the Session-Sender over UDP transport, and are used to collect the transmit and receive counters for the data traffic flow under measurement. The DLM reply probe packets are transmitted by the Session-Reflector to collect the transmit and receive counters for the data traffic flow under measurement in the reverse direction.

The DLM probe packets carry user-configured destination UDP port. The destination UDP port 862 is not used for the DLM probe packets. The user-configured destination UDP port follows the guidelines described in Section 4.1 of [RFC8762]. Different destination UDP port is used for DLM probe packets than the STAMP test packets defined in [RFC8762]. Hence, the Session-Sender and the Session-Reflector do not require backwards compatibility and support for STAMP.

A DLM session is identified by the 4-tuple (source and destination IP addresses, source and destination UDP port numbers). A DLM Session-Sender MAY generate a locally unique Sender Session Identifier (SSID). The SSID is a two-octet, non-zero unsigned integer. The SSID generation policy is implementation specific. An implementation MUST NOT assign the same identifier to different DLM sessions. A Session-Sender uses the SSID to identify a DLM session.

The DLM Session-Reflector operates in the Stateless mode.

In this document, the examples of DLM probe packets are shown with UDP header, however, the probe packets can be encapsulated with a different header based on the transport protocol used in the network.

4. Session-Sender Direct Loss Measurement Probe Packet

In this document, base Session-Sender DLM probe packet formats are defined as shown in Figure 1 and Figure 2 for unauthenticated and authenticated modes, respectively. They are stand-alone DLM probe packet formats to carry the counters for the data traffic flow under measurement. The DLM probe packet formats are similar to the base STAMP test packet formats (for example the locations of the Counters vs. STAMP Timestamps).

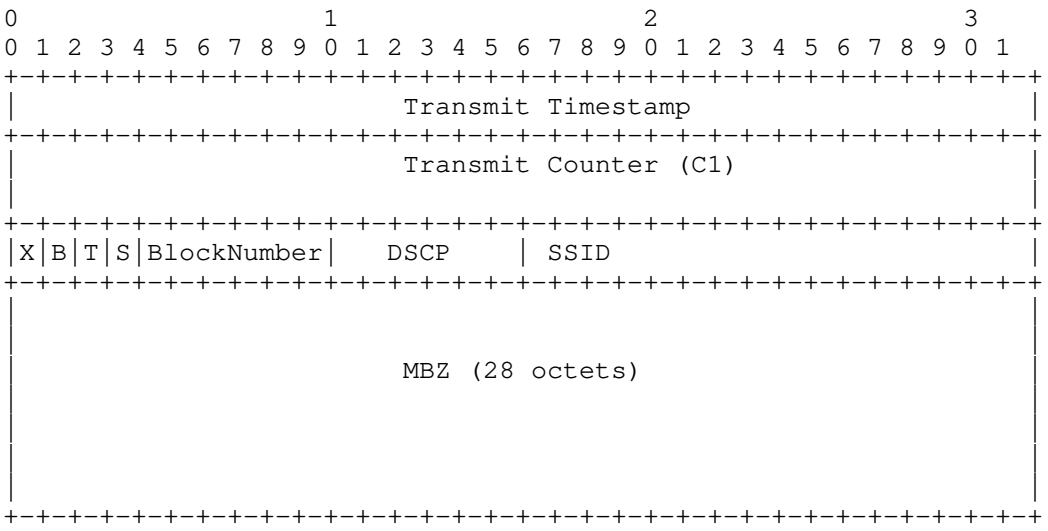

Figure 1: Session-Sender Direct Loss Measurement Probe Packet - Unauthenticated Mode

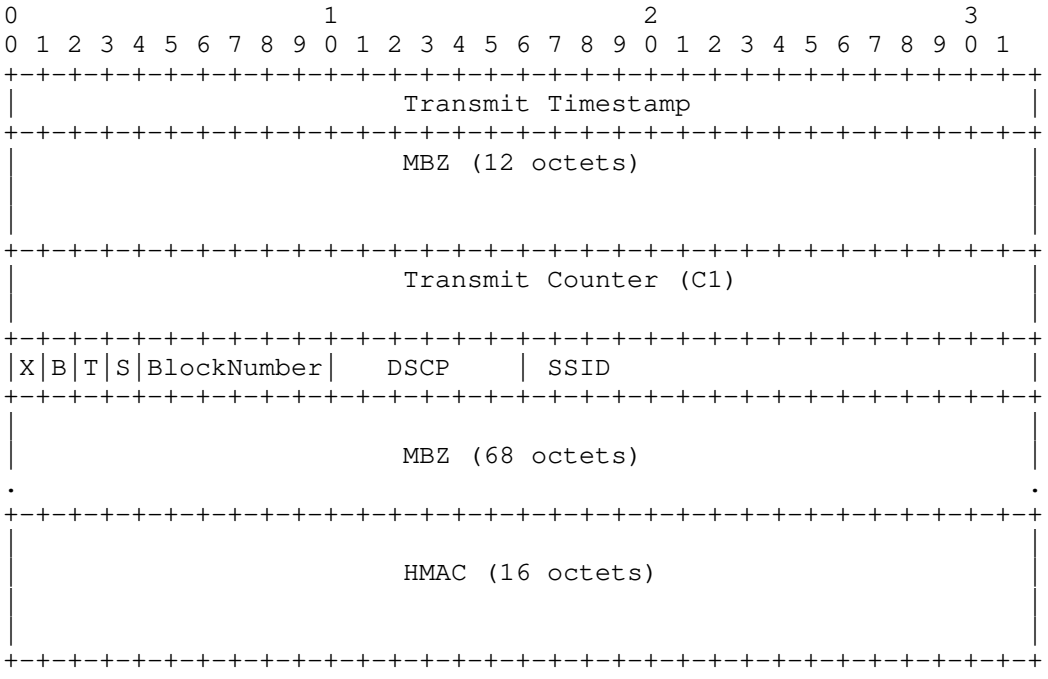

Figure 2: Session-Sender Direct Loss Measurement Probe Packet - Authenticated Mode

Fields are defined as the following:

Transmit Timestamp (32-bit): This is 32-bit nano-sec field of the PTPv2 timestamp on transmit side. This field may carry Sequence Number instead of PTPv2 timestamp.

Transmit Counter (64-bit): The number of packets or octets transmitted by the Session-Sender in the DLM probe packet. The counter is always written at the well-known fixed location in the DLM probe packet. This is an important property for hardware-based implementation (e.g., in an ASIC), e.g., for point-to-point links and circuits. Counter is for the data traffic flow under measurement.

XBTS Flags (3-bit): The meanings of the Flag bits are:

X: Extended counter format indicator. Indicates the use of extended (64-bit) counter values. Initialized to 1 upon creation (and prior to transmission) of a DLM probe packet. Set to 0 when the DLM probe packet is transmitted or received over an interface that writes 32-bit counter values.

B: Octet (byte) count. When set to 1, indicates that the Counter fields represent octet counts. The octet count applies to all packets within the DLM scope, and the octet count of a packet transmitted or received includes the total length of that packet (but excludes headers, labels, or framing of the channel itself). When set to 0, indicates that the Counter fields represent packet counts.

T: Traffic-class-specific measurement indicator. Set to 1 when the DLM session is scoped to data packets of a particular traffic class (DSCP value), and 0 otherwise. When set to 1, the DSCP field of the DLM probe packet indicates the measured traffic class.

S: Sequence Number indicator. When set to 1, it indicates that the Transmit Timestamp field contains Sequence Number (instead of PTPv2 timestamp).

DSCP (6-bit): DSCP of the data traffic flow being measured when T flag is set.

Block Number (6-bit): The Direct Loss Measurement using Alternate-Marking Method [RFC9341] requires collecting Block Number of the counters for the data traffic flow under measurement. To be able to correlate the transmit and receive counters of the matching Block Number, the Block Number of the counters carried in the DLM probe packets.

SSID (16-bit): DLM Sender Session Identifier.

HMAC: The use of the HMAC field is described in Section 4.4 of [RFC8762]. HMAC uses its own key and the mechanism to distribute the HMAC key is outside the scope of this document.

MBZ: Must be Zero. It MUST be all zeroed on the transmission and MUST be ignored on receipt.

5. Session-Reflector Direct Loss Measurement Probe Packet

The Session-Reflector receives the DLM Session-Sender probe packet and verifies it. If the DLM probe packet is validated, the Session-Reflector that supports this specification prepares and transmits the DLM reply probe packet. In this document, Session-Reflector DLM reply probe packet formats are defined as shown in Figure 3 and Figure 4, for unauthenticated and authenticated modes, respectively.

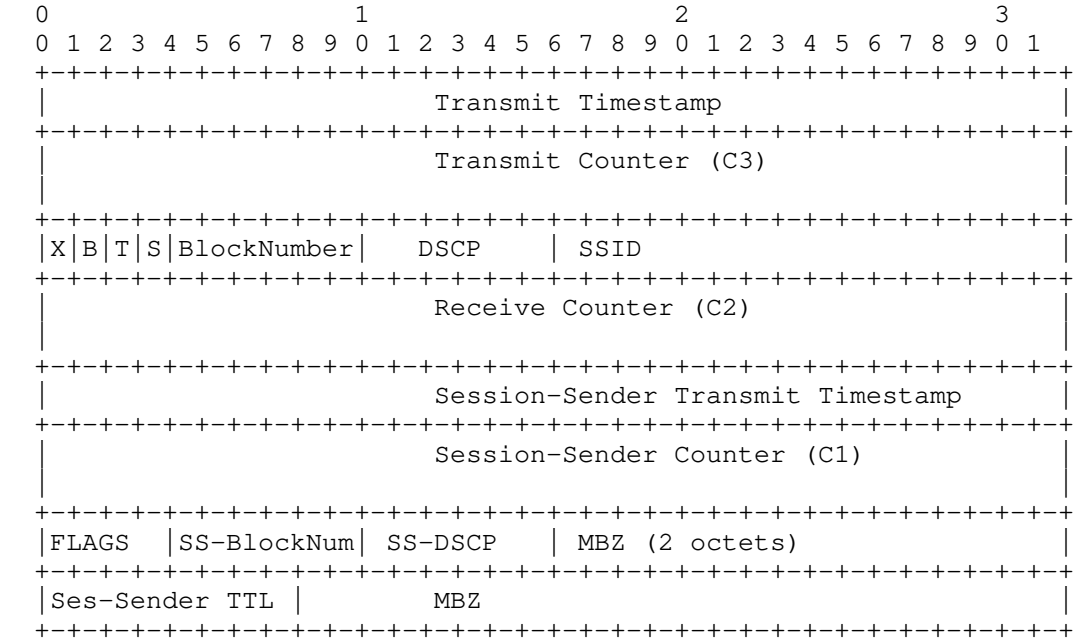

Figure 3: Session-Reflector Direct Loss Measurement Probe Packet - Unauthenticated Mode

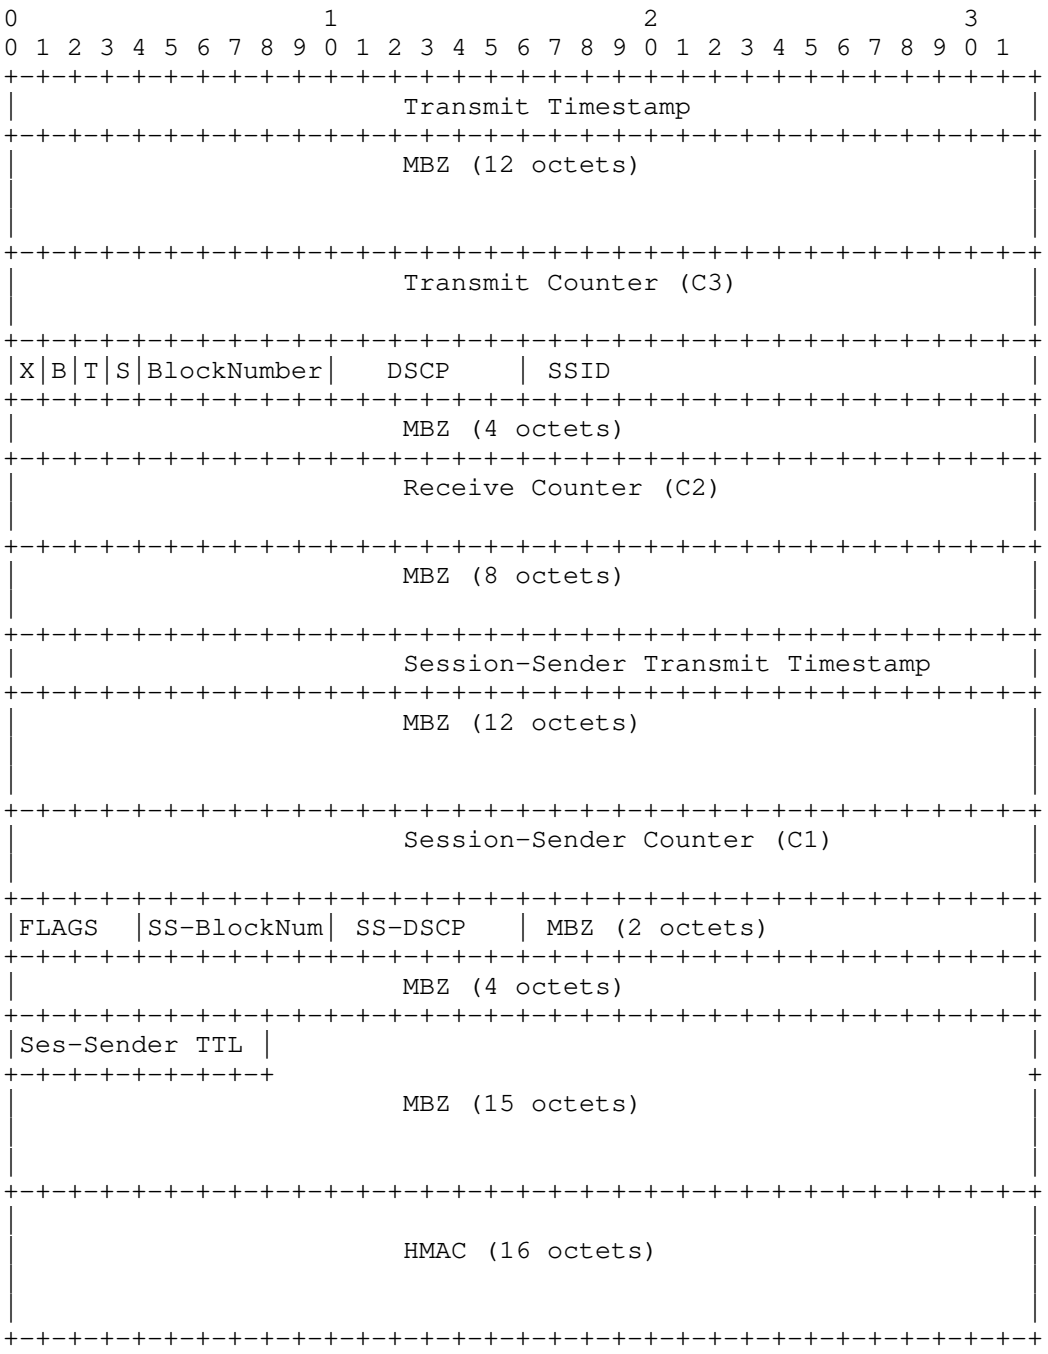
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
Transmit Timestamp
MBZ (12 octets)
Transmit Counter (C3)
X B T S BlockNumber DSCP SSID
MBZ (4 octets)
Receive Counter (C2)
MBZ (8 octets)
Session-Sender Transmit Timestamp
MBZ (12 octets)
Session-Sender Counter (C1)
FLAGS SS-BlockNum SS-DSCP MBZ (2 octets)
MBZ (4 octets)
Ses-Sender TTL
MBZ (15 octets)
HMAC (16 octets)

Figure 4: Session-Reflector Direct Loss Measurement Probe Packet - Authenticated Mode

Fields are defined as the following:

Transmit Timestamp (32-bit): This is 32-bit nano-sec field of the PTPv2 timestamp on transmit side. This field may carry Sequence Number instead of PTPv2 timestamp.

Transmit Counter (64-bit): The number of packets or octets transmitted by the Session-Reflector in the DLM reply probe packet. Counter is for the reverse direction data traffic flow under measurement. The Session-Reflector writes the Transmit Counter at the same location in the DLM reply probe packet as the Session-Sender DLM probe packet. This is an important property for hardware-based implementation (e.g., in an ASIC).

XBTS Flags (3-bit): The XBTS Flags for the reverse direction data traffic flow under measurement set using the same procedure defined for the Session-Sender DLM probe packet.

DSCP (6-bit): Set for the reverse direction data traffic flow under measurement using the same procedure defined for the Session-Sender DLM probe packet.

Block Number (6-bit): Set for the reverse direction data traffic flow under measurement using the same procedure defined for the Session-Sender DLM probe packet.

SSID: SSID is the exact copy of the SSID in the received Session-Sender DLM probe packet.

Receive Counter (64-bit): The number of packets or octets received at the Session-Reflector. It is written by the Session-Reflector in the DLM reply probe packet. Counter is for the data traffic flow under measurement.

Session-Sender Counter (64-bit): This is the exact copy of the Transmit Counter from the received Session-Sender DLM probe packet.

Session-Sender Transmit Timestamp (32-bit): This is the exact copy of the Transmit Timestamp from the received Session-Sender DLM probe packet.

Session-Sender Block Number: This is the exact copy of the Block Number from the received Session-Sender DLM probe packet.

Session-Sender FLAGS: This is the exact copy of the XBTS Flags from the received Session-Sender DLM probe packet.

Session-Sender DSCP: This is the exact copy of the DSCP from the received Session-Sender DLM probe packet.

Session-Sender TTL: The Session-Sender TTL field is one octet long, and its value is the copy of the TTL field in IPv4 (or Hop Limit in IPv6) from the received Session-Sender DLM probe packet.

6. Data Loss Calculation

Using the Counters C1, C2, C3 and C4 as per reference topology, from the nth and (n-1)th DLM probe packets, packet loss and byte loss for the data traffic flow can be calculated as follows:

Transmit Loss TxL[n-1, n] = (C1[n] - C1[n-1]) - (C2[n] - C2[n-1])

Receive Loss RxL[n-1, n] = (C3[n] - C3[n-1]) - (C4[n] - C4[n-1])

The Total Transmit and Receive Loss are calculated as follows:

Total Transmit Loss = TxL[1, 2] + TxL[2, 3] + ...

Total Receive Loss = RxL[1, 2] + RxL[2, 3] + ...

These values are updated each time a DLM reply probe packet is received and processed at the Session-Sender, and they represent the Total Transmit and Total Receive Loss since the DLM session was initiated. When computing the values TxL[n-1,n] and RxL[n-1,n], the possibility of counter wrap must be taken into account.

When using Alternate-Marking Method, all Counters used for loss calculation belongs to the same Block Number, as described in Section 3.1 of [RFC9341].

7. Optional Extensions

There are currently no optional (TLV) extensions defined for the DLM probe packets.

8. Integrity Protection and Confidentiality Protection

The integrity protection and confidentiality protection specified in [RFC8762] also apply to the procedures defined in this document.

9. Operational Considerations

The operational considerations specified in [RFC8762] also apply to the procedures defined in this document.

10. Security Considerations

The DLM protocol defined in this document is intended for deployment in a single operator network domain. As such, the Session-Sender address and Session-Reflector address are provisioned by the operator for the DLM session. It is assumed that the operator has verified the integrity of the path and identity of the far-end Session-Reflector.

If desired, attacks can be mitigated by performing basic validation and sanity checks, at the Session-Sender, of the counter fields (e.g., packet loss is not negative) in received reply probe packets. The minimal state associated with these protocols also limits the extent of measurement disruption that can be caused by a corrupt or invalid probe packet to a single probe cycle.

The DLM protocol uses UDP port that could become a target of denial of service (DoS) or could be used to aid on-path attacks. Thus, the security considerations and measures to mitigate the risk of the attack documented in Section 6 of [RFC8545] equally apply to the STAMP extensions in this document.

The security considerations specified in [RFC8762] and [RFC8972] also apply to the protocol defined in this document. Specifically, the message integrity protection using HMAC, as defined in [RFC8762] Section 4.4, also apply to the procedure described in this document.

11. IANA Considerations

This document has no IANA actions.

12. References

12.1. 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>.

[RFC9341] Fioccola, G., Ed., Cociglio, M., Mirsky, G., Mizrahi, T., and T. Zhou, "Alternate-Marking Method", RFC 9341, DOI 10.17487/RFC9341, December 2022, <https://www.rfc-editor.org/info/rfc9341>.

[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>.

12.2. Informative References

[RFC7680] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Loss Metric for IP Performance Metrics (IPPM)", STD 82, RFC 7680, DOI 10.17487/RFC7680, January 2016, <https://www.rfc-editor.org/info/rfc7680>.

[RFC8545] Morton, A., Ed. and G. Mirsky, Ed., "Well-Known Port Assignments for the One-Way Active Measurement Protocol (OWAMP) and the Two-Way Active Measurement Protocol (TWAMP)", RFC 8545, DOI 10.17487/RFC8545, March 2019, <https://www.rfc-editor.org/info/rfc8545>.

[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>.

[ITU-Y1731] Recommendation ITU-TG.8013/Y.1731: https://www.itu.int/rec/T-REC-G.8013-201508-I/en, "G.8013/Y.1731 : Operations, administration and maintenance (OAM) functions and mechanisms for Ethernet-based networks", August 2015.

[SRV6-PM-TNSM] Loreti, P., Mayer, A., Lungaroni, P., Lombardo, F., Scarpitta, C., Sidoretti, G., Bracciale, L., Ferrari, M., Salsano, S., Abdelsalam, A., Gandhi, R., and C. Filsfils, IEEE Transactions on Network and Service Management, "SRv6-PM: Performance Monitoring of SRv6 Networks with a Cloud-Native Architecture: https://arxiv.org/pdf/2007.08633.pdf", February 2021.

[SRV6-PM-IEEE] Loreti, P., Mayer, A., Lungaroni, P., Salsano, S., Gandhi, R., and C. Filsfils, IEEE International Conference on High

Performance Switching and Routing, "Implementation of Accurate Per-Flow Packet Loss Monitoring in Segment Routing over IPv6 Networks: https://arxiv.org/pdf/2004.11414.pdf", May 2020.

Acknowledgments

The authors would like to thank Greg Mirsky, Tianran Zhou, Gyan Mishra, Zhenqiang Li, Reshad Rahman, Cheng Li, and Yali Wang for the comments on Direct Loss Measurement. The authors would like to thank Pierpaolo Loreti, Stefano Salsano, and the team for the Open Source implementation of SRv6-PM Loss Monitoring and its publications in [SRV6-PM-TNSM] and [SRV6-PM-IEEE]. The authors would like to acknowledge the earlier work on the loss measurement using TWAMP described in draft-xiao-ippm-twamp-ext-direct-loss.

Authors' Addresses

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

Clarence Filsfils
Cisco Systems, Inc.
Email: cfilsfil@cisco.com

Daniel Voyer
Bell Canada
Email: daniel.voyer@bell.ca

Mach(Guoyi) Chen
Huawei
Email: mach.chen@huawei.com

Bart Janssens
Colt
Email: Bart.Janssens@colt.net