

Workgroup: IPPM Working Group
Internet-Draft: draft-ietf-ippm-stamp-srpm-05
Published: 11 August 2022
Intended Status: Standards Track
Expires: 12 February 2023
Authors: R. Gandhi, Ed. C. Filsfils
 Cisco Systems, Inc. Cisco Systems, Inc.
 D. Voyer M. Chen B. Janssens R. Foote
 Bell Canada Huawei Colt Nokia
Simple TWAMP (STAMP) Extensions for Segment Routing Networks

Abstract

Segment Routing (SR) leverages the source routing paradigm. SR is applicable to both Multiprotocol Label Switching (SR-MPLS) and IPv6 (SRv6) forwarding planes. This document specifies RFC 8762 (Simple Two-Way Active Measurement Protocol (STAMP)) extensions for SR networks, for both SR-MPLS and SRv6 forwarding planes by augmenting the optional extensions defined in RFC 8972.

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 12 February 2023.

Copyright Notice

Copyright (c) 2022 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 Revised BSD License text as described in

Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Conventions Used in This Document](#)
 - [2.1. Requirements Language](#)
 - [2.2. Abbreviations](#)
 - [2.3. Reference Topology](#)
- [3. Verification Check Flag in TLV](#)
- [4. Destination Node Address TLV](#)
- [5. Return Path TLV](#)
 - [5.1. Return Path Sub-TLVs](#)
 - [5.1.1. Return Path Control Code Sub-TLV](#)
 - [5.1.2. Return Address Sub-TLV](#)
 - [5.1.3. Return Segment List Sub-TLVs](#)
- [6. Security Considerations](#)
- [7. IANA Considerations](#)
- [8. References](#)
 - [8.1. Normative References](#)
 - [8.2. Informative References](#)
- [Acknowledgments](#)
- [Authors' Addresses](#)

1. Introduction

Segment Routing (SR) leverages the source routing paradigm for Software Defined Networks (SDNs). SR is applicable to both Multiprotocol Label Switching (SR-MPLS) and IPv6 (SRv6) forwarding planes [[RFC8402](#)]. SR Policies as defined in [[RFC9256](#)] are used to steer traffic through a specific, user-defined paths using a stack of Segments. A comprehensive SR Performance Measurement (PM) toolset is one of the essential requirements to measure network performance to provide Service Level Agreements (SLAs).

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. [[RFC8972](#)] defines optional extensions, in the form of TLVs, for STAMP. Note that the YANG data model defined in [[I-D.ietf-ippm-stamp-yang](#)] can be used to provision the STAMP Session-Sender and STAMP Session-Reflector.

The STAMP test packets are transmitted along an IP path between a Session-Sender and a Session-Reflector to measure performance delay and packet loss along that IP path. It may be desired in SR networks that the same path (same set of links and nodes) between the Session-Sender and Session-Reflector is used for the STAMP test

packets in both directions. This is achieved by using the STAMP [[RFC8762](#)] extensions for SR-MPLS and SRv6 networks specified in this document by augmenting the optional extensions defined in [[RFC8972](#)].

2. Conventions Used in This Document

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD 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.

2.2. Abbreviations

MPLS: Multiprotocol Label Switching.

PM: Performance Measurement.

SID: Segment ID.

SL: Segment List.

SR: Segment Routing.

SR-MPLS: Segment Routing with MPLS forwarding plane.

SRv6: Segment Routing with IPv6 forwarding plane.

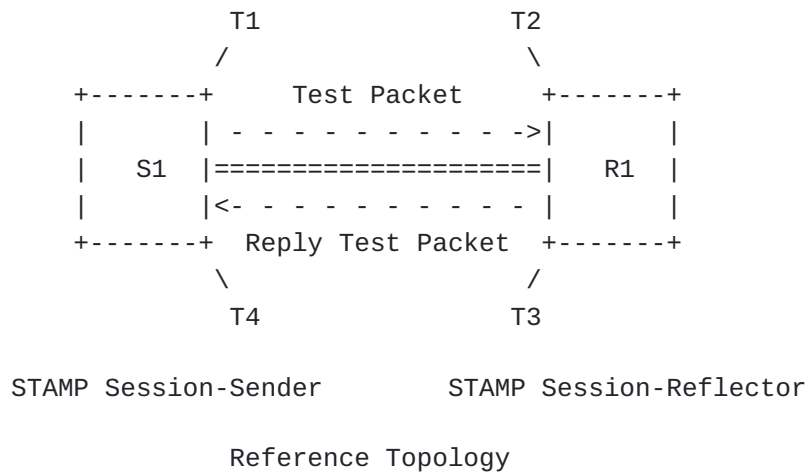
SSID: STAMP Session Identifier.

STAMP: Simple Two-Way Active Measurement Protocol.

2.3. Reference Topology

In the reference topology shown below, the STAMP Session-Sender S1 initiates a STAMP test packet and the STAMP Session-Reflector R1 transmits a reply STAMP test packet. The reply test packet may be transmitted to the Session-Sender S1 on the same path (same set of links and nodes) or a different path in the reverse direction from the path taken towards the Session-Reflector R1.

The nodes S1 and R1 may be connected via a link or an SR path [[RFC8402](#)]. The link may be a physical interface, virtual link, or Link Aggregation Group (LAG) [[IEEE802.1AX](#)], or LAG member. The SR path may be an SR Policy [[RFC9256](#)] on node S1 (called head-end) with destination to node R1 (called tail-end).



3. Verification Check Flag in TLV

The STAMP TLV option in [\[RFC8972\]](#) defines the use of the 8-bit flags field common to all STAMP TLVs.

A one-bit flag called Verification Check (V) is defined at bit position (3) in the flags field of the STAMP TLV. A Session-Sender MUST set the V flag to 0 before transmitting an extended STAMP test packet when reply test packet is required. A Session-Reflector MUST set the V flag to 1 for any STAMP TLV that it supports that includes an instruction that cannot be met or conflicts with the Session-Reflector processing or capability. The V flag MUST be set to 0 by the Session-Reflector when the instruction from the request is met.

A Session-Sender MUST set the V flag to 1 before transmitting an extended STAMP test packet when test packet reply is not required. A Session-Reflector MUST NOT reply and MUST drop the test packet if the Session-Reflector determined for any STAMP TLV that it supports that includes an instruction that cannot be met or is in conflict with the Session-Reflector policy or capability.

For the STAMP Session-Reflector that supports the STAMP TLV extensions defined in this document, the test packets may carry additional instructions in a TLV for the Session-Reflector to follow. In this case, the V flag provides feedback to the Session-Sender if the Session-Reflector was able to follow that instruction. For example, Session-Reflector supports the TLV and it is well-formed, the STAMP test packet including all the TLVs was successfully processed but the additional instruction in one of the TLVs was not followed. The Session-Reflector can either drop the packet or return the packet with the error back to the Session-Sender, based on the request in the test packet and local policy.

4. Destination Node Address TLV

The Session-Sender may need to transmit test packets to the Session-Reflector with a different destination address that is not matching an address on the Session-Reflector e.g. when the STAMP test packet is encapsulated by a tunneling protocol or an MPLS Segment List with destination IPv4 address from 127/8 range or Segment Routing Header (SRH) with destination IPv6 address ::1/128. When using IPv4 destination address from 127/8 range (e.g. for testing in an ECMP environment), the STAMP test packet may not reach the intended Session-Reflector in an error condition, and an un-intended node may transmit reply test packet resulting in reporting of invalid measurement metrics.

[RFC8972] defines STAMP test packets that can include one or more optional TLVs. In this document, the TLV type (value 9) is defined for the Destination Node Address TLV for the STAMP test packet [RFC8972]. The format of the Destination Node Address TLV is shown in Figure 1:




Figure 1: Destination Node Address TLV Format

The Length field is used to decide the Address Family of the Address.

The STAMP TLV Flags are set using the procedures described in [RFC8972].

The Destination Node Address TLV is optional. The Destination Node Address TLV indicates the address of the intended Session-Reflector node of the test packet that is also used to uniquely identify the STAMP session on the Session-Reflector when the optional SSID is not sent. The Session-Reflector MUST add the received Destination Node Address TLV in the reply test packet to ensure the symmetric reply test packet size and to transmit the STAMP TLV Flags to the Session-Sender.

A Session-Sender MUST set the V flag to 0 before transmitting an extended STAMP test packet when test packet reply is required. A Session-Reflector that supports this TLV, MUST set the V flag in the reply test packet to 1 if the Session-Reflector determined that it

is not the intended Destination as identified in the Destination Node Address TLV. Otherwise, the Session-Reflector MUST set the V flag in the reply test packet to 0.

A Session-Sender MUST set the V flag to 1 before transmitting an extended STAMP test packet when test packet reply is not required. A Session-Reflector that supports this TLV, MUST NOT reply and MUST drop the test packet if the Session-Reflector determined that it is not the intended Destination as identified in the Destination Node Address TLV.

5. Return Path TLV

For end-to-end SR paths, the Session-Reflector may need to transmit the reply test packet on a specific return path. The Session-Sender can request this in the test packet to the Session-Reflector using a Return Path TLV. With this TLV carried in the Session-Sender test packet, signaling and maintaining dynamic SR network state for the STAMP sessions on the Session-Reflector are avoided.

There are two modes defined for the behaviors on the Session-Reflector in Section 4 of [\[RFC8762\]](#). A Stateful Session-Reflector that requires configuration that must match all Session-Sender parameters, including Source Address, Destination Address, Source UDP Port, Destination UDP Port, and possibly SSID (assuming the SSID is configurable and not auto-generated). In this case, a local policy can be used to direct the test packet by creating additional states for the STAMP sessions on the Session-Reflector. In the case of promiscuous operation, the Stateless Session-Reflector will require an indication of how to return the test packet on a specific path, for example, in an ECMP environment.

For links, the Session-Reflector may need to transmit the reply test packet on the same incoming link in the reverse direction. The Session-Sender can request this in the test packet to the Session-Reflector using a Return Path TLV.

[\[RFC8972\]](#) defines STAMP test packets that can include one or more optional TLVs. In this document, the TLV Type (value 10) is defined for the Return Path TLV that carries the return path for the Session-Sender test packet. The format of the Return Path TLV is shown in Figure 2:

When Return Path Sub-TLV is present in the Session-Sender test packet, the Session-Reflector that supports this TLV, MUST transmit reply test packet using the return path information specified in the Return Path Sub-TLV.

A Return Path TLV MUST NOT contain both Control Code Sub-TLV as well as Return Address or Return Segment List Sub-TLV.

5.1.1. Return Path Control Code Sub-TLV

The format of the Return Path Control Code Sub-TLV is shown in Figure 3. The Type of the Return Path Control Code Sub-TLV is defined as following:

*Type (value 1): Return Path Control Code. The Session-Sender can request the Session-Reflector to transmit the reply test packet based on the flags defined in the Control Code field.




Figure 3: Control Code Sub-TLV in Return Path TLV

Control Code Flags (32-bit): Defined as follows.

0x0: No Reply Requested.

0x1: Reply Requested on the Same Link.

When Control Code flag is set to 0x0 in the Session-Sender test packet, the Session-Reflector does not transmit reply test packet to the Session-Sender and terminates the STAMP test packet. Only the one-way measurement is applicable in this case. Optionally, the Session-Reflector may locally stream performance metrics via telemetry using the information from the received test packet. All other Return Path Sub-TLVs MUST be ignored in this case.

When Control Code flag is set to 0x1 in the Session-Sender test packet, the Session-Reflector transmits the reply test packet over the same incoming link where the test packet is received in the reverse direction towards the Session-Sender. All other Return Path Sub-TLVs MUST be ignored in this case.

5.1.2. Return Address Sub-TLV

The STAMP reply test packet may be transmitted to the Session-Sender to a different destination address on the Session-Sender using Return Path TLV. For this, the Session-Sender can specify in the test packet the receiving destination node address for the Session-Reflector reply test packet. When transmitting the STAMP test packet to a different destination address, the Session-Sender MUST follow the procedure defined in Section 4.3 of [[RFC8762](#)].

The format of the Return Address Sub-TLV is shown in Figure 4. The Address Family field indicates the type of the address, and it SHALL be set to one of the assigned values in the "IANA Address Family Numbers" registry. The Type of the Return Address Sub-TLV is defined as following:

*Type (value 2): Return Address. Destination node address of the Session-Reflector reply test packet different than the Source Address in the Session-Sender test packet.

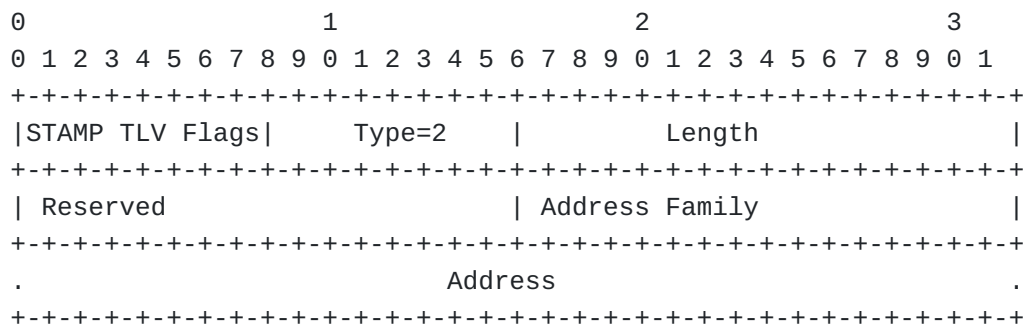


Figure 4: Return Address Sub-TLV in Return Path TLV

5.1.3. Return Segment List Sub-TLVs

The format of the Segment List Sub-TLVs in the Return Path TLV is shown in Figures 5, 6, and 7. The segment entries MUST be in network order. The Segment List Sub-TLV can be one of the following Types:

*Type (value 3): SR-MPLS Label Stack of the Return Path

*Type (value 4): SRv6 Segment List of the Return Path

*Type (value 5): Structured SRv6 Segment List of the Return Path

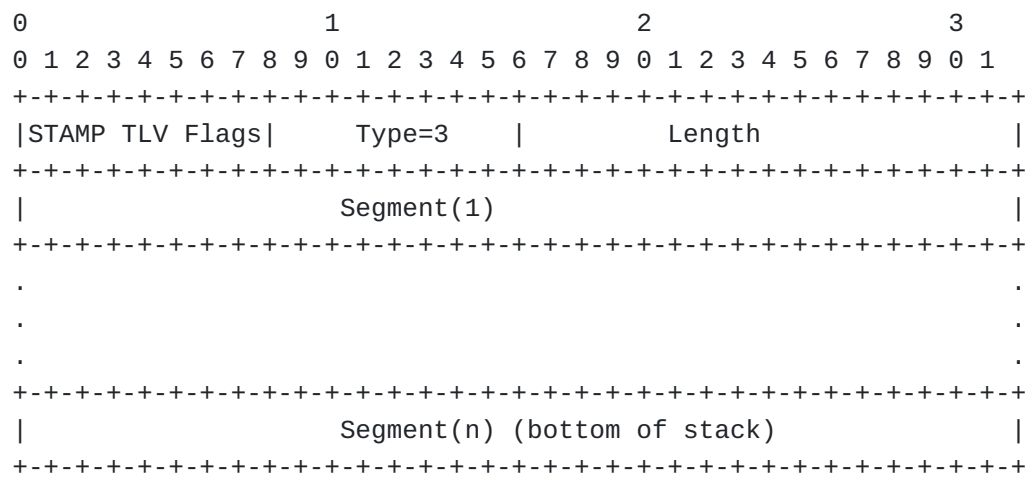


Figure 5: SR-MPLS Segment List Sub-TLV in Return Path TLV

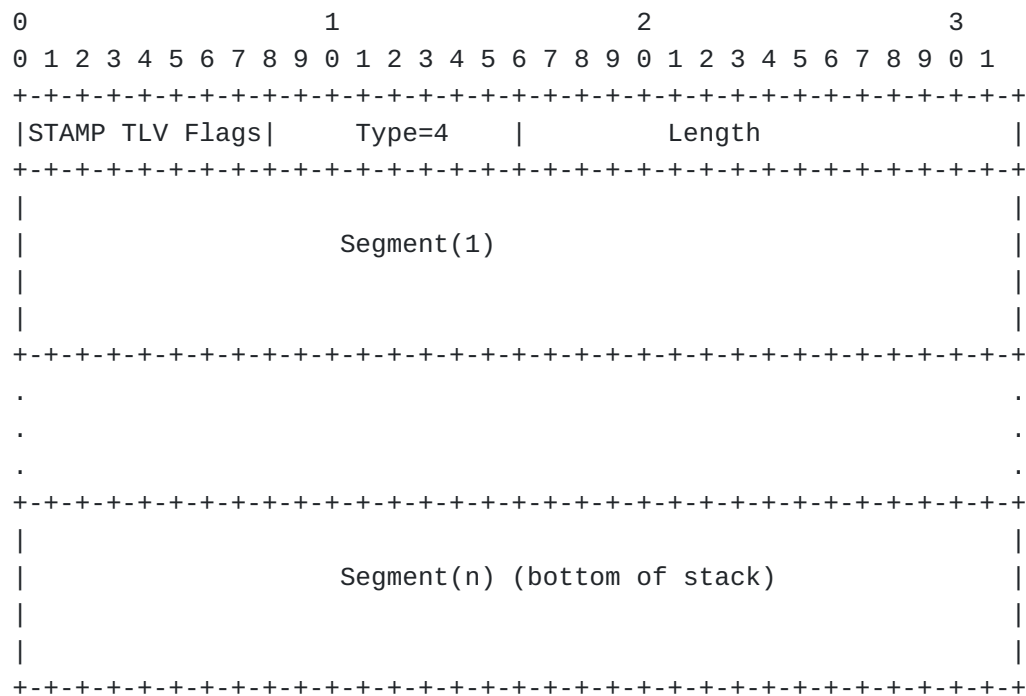


Figure 6: SRv6 Segment List Sub-TLV in Return Path TLV

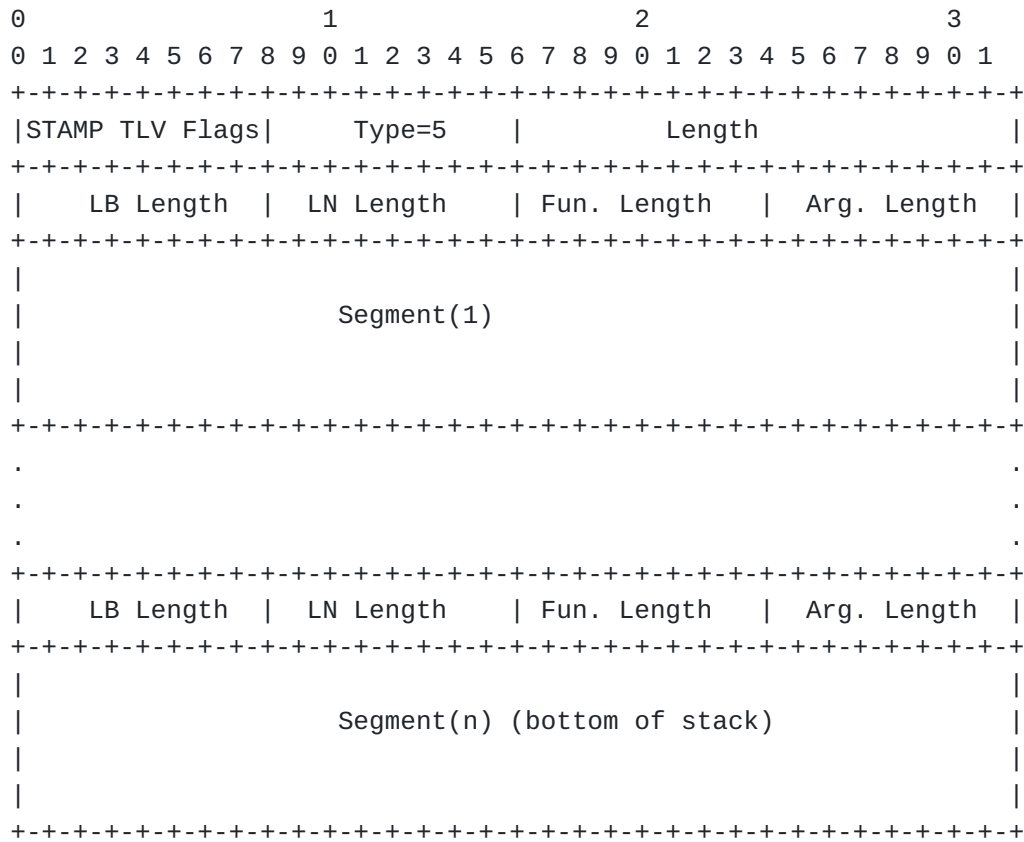


Figure 7: Structured SRv6 Segment List Sub-TLV in Return Path TLV

The SR-MPLS Label Stack contains a list of 32-bit Label Stack Entry (LSE) that includes a 20-bit label value, 8-bit TTL value, 3-bit TC value and 1-bit EOS field. An SR-MPLS Label Stack Sub-TLV may carry only Binding SID Label [[I-D.ietf-pce-binding-label-sid](#)] of the Return SR-MPLS Policy.

An SRv6 Segment List Sub-TLV and Structured SRv6 Segment List Sub-TLV may carry only Binding SID [[I-D.ietf-pce-binding-label-sid](#)] of the Return SRv6 Policy.

A Structured SRv6 Segment List Sub-TLV is used carry the structure and behavior for SRv6 SIDs [[RFC8986](#)] used in the Return SRv6 path as shown in Figure 7. The structure is intended for informational use by the control and management planes. The fields in the structure of the Sub-TLV are defined as follows [[RFC8986](#)]:

*LB Length: 1 octet. SRv6 SID Locator Block (LB) length in bits.

*LN Length: 1 octet. SRv6 SID Locator Node (LN) length in bits.

*Fun. Length: 1 octet. SRv6 SID Function length in bits.

*Arg. Length: 1 octet. SRv6 SID Arguments length in bits.

In Structured SRv6 Segment List Sub-TLV, the sum of all four sizes MUST be less than or equal to 128 bits. If the sum of all four sizes is larger than 128 bits, the Sub-TLV MUST NOT be used by the Session-Reflector.

The Session-Sender MUST only insert one Segment List Return Path Sub-TLV in the test packet. The Session-Reflector MUST only process the first Segment List Return Path Sub-TLV in the test packet and ignore other Segment List Return Path Sub-TLVs if present.

Note that in addition to P2P SR paths, the Return Segment List Sub-TLV is also applicable to P2MP SR paths. For example, for P2MP SR paths, it may only carry the Node Segment Identifier of the Session-Sender in order for the reply test packet to follow an SR path to the Session-Sender.

6. Security Considerations

The usage of STAMP protocol is intended for deployment in limited domains [[RFC8799](#)]. As such, it assumes that a node involved in STAMP protocol operation has previously verified the integrity of the path and the 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 timestamp fields in received reply test packets. The minimal state associated with these protocols also limits the extent of measurement disruption that can be caused by a corrupt or invalid test packet to a single test cycle.

The security considerations specified in [[RFC8762](#)] and [[RFC8972](#)] also apply to the extensions 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.

STAMP uses the well-known UDP port number that could become a target of denial of service (DoS) or could be used to aid man-in-the-middle (MITM) 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 STAMP extensions defined in this document may be used for potential "proxying" attacks. For example, a Session-Sender may specify a return path that has a destination different from that of the Session-Sender. But normally, such attacks will not happen in an SR domain where the Session-Senders and Session-Reflectors belong to the same domain. In order to prevent using the extension defined in this document for proxying any possible attacks, the return path has destination to the same node where the forward path is from. The Session-Reflector may drop the Session-Sender test packet when it

cannot determine whether the Return Path has the destination to the Session-Sender. That means, the Session-Sender should choose a proper source address according to the specified Return Path to help the Session-Reflector to make that decision.

7. IANA Considerations

IANA has created the "STAMP TLV Types" registry for [[RFC8972](#)]. IANA has early allocated a value for the Destination Address TLV Type and a value for the Return Path TLV Type from the IETF Review TLV range of the same registry.

Value	Description	Reference
9 (Early Allocation)	Destination Node Address	This document
10 (Early Allocation)	Return Path	This document

Table 1: STAMP TLV Types

IANA is requested to create a sub-registry for "Return Path Sub-TLV Type". All code points in the range 1 through 175 in this registry shall be allocated according to the "IETF Review" procedure as specified in [[RFC8126](#)]. Code points in the range 176 through 239 in this registry shall be allocated according to the "First Come First Served" procedure as specified in [[RFC8126](#)]. Remaining code points are allocated according to [Table 2](#):

Value	Description	Reference
1 - 175	IETF Review	This document
176 - 239	First Come First Served	This document
240 - 251	Experimental Use	This document
252 - 254	Private Use	This document

Table 2: Return Path Sub-TLV Type Registry

IANA is requested to allocate the values for the following Sub-TLV Types from this registry.

Type	Description	Reference
0	Reserved	This document
1	Return Path Control Code	This document
2	Return Address	This document
3	SR-MPLS Label Stack of the Return Path	This document
4	SRv6 Segment List of the Return Path	This document
5	Structured SRv6 Segment List of the Return Path	This document
255	Reserved	This document

Table 3: Return Path Sub-TLV Types

IANA has created the "STAMP TLV Flags" subregistry. IANA has early allocated the following bit position in the "STAMP TLV Flags" subregistry.

Bit Position	Symbol	Description	Reference
3 (Early Allocation)	V	Verification Check	This document

Table 4: STAMP TLV Flags

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>>.
- [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>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

8.2. Informative References

- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.

[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>>.

[RFC8799]

Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <<https://www.rfc-editor.org/info/rfc8799>>.

[RFC9256]

Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.

[I-D.ietf-pce-binding-label-sid]

Sivabalan, S., Filsfils, C., Tantsura, J., Previdi, S., and C. L. (editor), "Carrying Binding Label/Segment Identifier in PCE-based Networks.", Work in Progress, Internet-Draft, draft-ietf-pce-binding-label-sid-15, 20 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-pce-binding-label-sid-15.txt>>.

[I-D.ietf-ippm-stamp-yang]

Mirsky, G., Min, X., and W. S. Luo, "Simple Two-way Active Measurement Protocol (STAMP) Data Model", Work in Progress, Internet-Draft, draft-ietf-ippm-stamp-yang-10, 10 July 2022, <<https://www.ietf.org/archive/id/draft-ietf-ippm-stamp-yang-10.txt>>.

[IEEE802.1AX]

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

Acknowledgments

The authors would like to thank Thierry Couture for the discussions on the use-cases for Performance Measurement in Segment Routing. The authors would also like to thank Greg Mirsky, Mike Koldychev, Gyan Mishra, Tianran Zhou, Al Mortons, Reshad Rahman, Zhenqiang Li, Frank Brockners, and Cheng Li for providing comments and suggestions.

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

Richard Foote
Nokia

Email: footer.foote@nokia.com