

IPPM Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 14, 2021

R. Gandhi, Ed.
C. Filsfils
Cisco Systems, Inc.
D. Voyer
Bell Canada
M. Chen
Huawei
B. Janssens
Colt
February 10, 2021

Simple TWAMP (STAMP) Extensions for Segment Routing Networks
draft-gandhi-ippm-stamp-srpm-02

Abstract

Segment Routing (SR) leverages the source routing paradigm. SR is applicable to both Multiprotocol Label Switching (SR-MPLS) and IPv6 (SRv6) data planes. This document specifies [RFC 8762](#) (Simple Two-Way Active Measurement Protocol (STAMP)) extensions for SR networks, for both SR-MPLS and SRv6 data planes.

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 14, 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.	Conventions Used in This Document	3
2.1.	Requirements Language	3
2.2.	Abbreviations	3
2.3.	Reference Topology	4
3.	Destination Node Address TLV	4
4.	Return Path TLV	5
4.1.	Return Path Sub-TLVs	6
4.1.1.	Return Path Control Code Sub-TLV	6
4.1.2.	Return Address Sub-TLV	7
4.1.3.	Return Segment List Sub-TLVs	8
5.	Security Considerations	9
6.	IANA Considerations	9
7.	References	10
7.1.	Normative References	10
7.2.	Informative References	11
	Acknowledgments	11
	Authors' Addresses	11

[1.](#) Introduction

Segment Routing (SR) leverages the source routing paradigm and greatly simplifies network operations for Software Defined Networks (SDNs). SR is applicable to both Multiprotocol Label Switching (SR-MPLS) and IPv6 (SRv6) data planes [[RFC8402](#)]. SR Policies as defined in [[I-D.ietf-spring-segment-routing-policy](#)] are used to steer traffic through a specific, user-defined paths using a stack of Segments. Built-in SR Performance Measurement (PM) is one of the essential requirements 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](#)]. It eliminates the need for control protocol by using configuration and management model to provision and manage test sessions. [[RFC8972](#)] defines optional extensions for STAMP.

The STAMP supports two modes of STAMP Session-Reflector: Stateless and Stateful as described in [Section 4 of \[\[RFC8762\]\(#\)\]](#). In Stateless mode, maintenance of each STAMP test session on Session-Reflector is

avoided. In SR networks, as the state is in the packet, the signaling of the parameters and creating extra states in the network are undesired. Hence, Stateless mode of Session-Reflector is preferred in SR networks.

For performance delay and packet loss measurement, STAMP Session-Sender test packets are transmitted in-band on the same path as the data traffic flow under measurement to measure the delay and packet loss experienced by the data traffic flow. It is also desired in SR networks that the Session-Reflector reply test packets are transmitted in-band on the same path in the reverse direction. This is achieved by using the STAMP extensions defined in this document.

This document specifies [RFC 8762](#) (Simple Two-Way Active Measurement Protocol (STAMP)) extensions for SR networks, for both SR-MPLS and SRv6 data planes.

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 data plane.

SRv6: Segment Routing with IPv6 data plane.

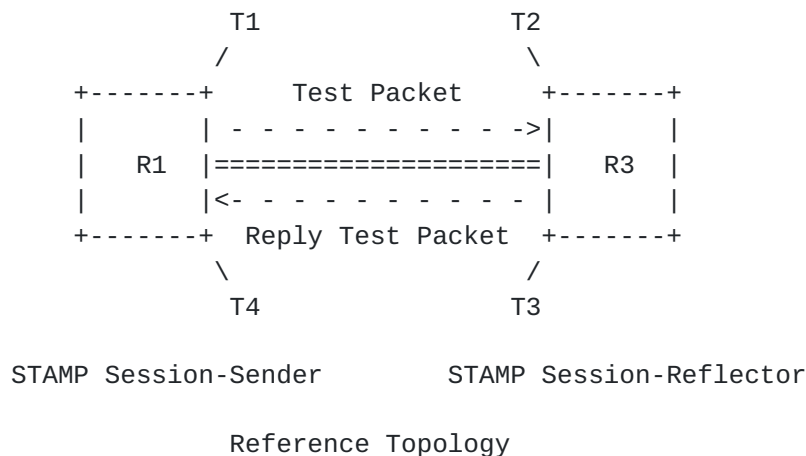
SSID: STAMP Session Identifier.

STAMP: Simple Two-way Active Measurement Protocol.

2.3. Reference Topology

In the reference topology shown below, the STAMP Session-Sender R1 initiates a STAMP test packet and the STAMP Session-Reflector R3 transmits a reply test packet. The reply test packet is transmitted back to the STAMP Session-Sender R1 on the same path or a different path in the reverse direction.

The nodes R1 and R3 may be connected via a link or there exists an SR path [RFC8402]. The link may be a physical interface, virtual link, or Link Aggregation Group (LAG) [IEEE802.1AX], or LAG member link. The SR path may be an SR Policy [I-D.ietf-spring-segment-routing-policy] on node R1 (called head-end) with destination to node R3 (called tail-end).



3. Destination Node Address TLV

The STAMP Session-Sender may need to transmit test packets to the STAMP Session-Reflector with a different destination address (for example IPv4 address from 127/8 range). In an error condition, the STAMP test packet may not reach the intended STAMP Session-Reflector, an un-intended node may transmit reply test packets resulting in reporting of invalid measurement metrics.

[RFC8972] defines STAMP test packets that can include one or more optional TLVs. In this document, Destination Node Address TLV (Type TBA1) is defined for STAMP test packet [RFC8972] and has the following format shown in Figure 1:

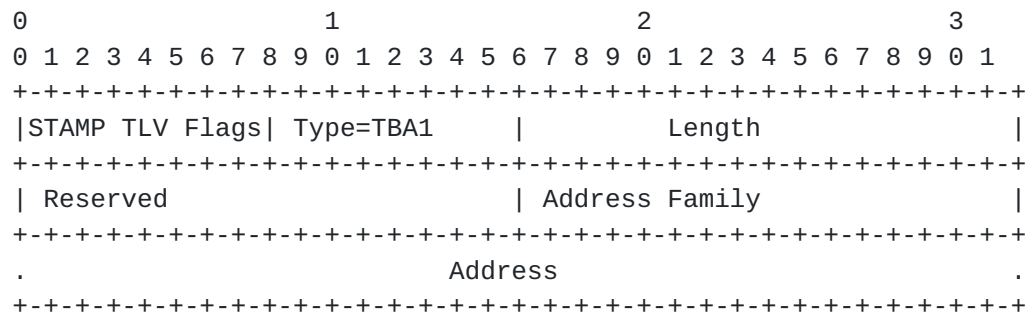


Figure 1: Destination Node Address TLV Format

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 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 destination node of the test packet. The STAMP Session-Reflector that supports this TLV, MUST NOT transmit reply test packet if it is not the intended destination node of the received test packet.

4. Return Path TLV

For end-to-end SR paths, the STAMP Session-Reflector may need to transmit the reply test packet on a specific return path. The STAMP Session-Sender can request this in the test packet to the STAMP Session-Reflector using a Return Path TLV. With this TLV carried in the STAMP Session-Sender test packet, the STAMP Session-Reflector (Stateless mode) does not require signaling and maintaining any additional dynamic state for the STAMP sessions for the end-to-end SR paths.

For links, the STAMP Session-Reflector may need to transmit the reply test packet on the same incoming link in the reverse direction. The STAMP Session-Sender can request this in the test packet to the STAMP Session-Reflector using a Return Path TLV. With this TLV carried in the STAMP Session-Sender test packet, the STAMP Session-Reflector (Stateless mode) does not require maintenance of any additional state for the STAMP sessions for the links.

[RFC8972] defines STAMP test packets that can include one or more optional TLVs. In this document, the TLV Type (value TBA2) is defined for the Return Path TLV that carries the return path for the

STAMP Session-Sender test packet. The format of the Return Path TLV is shown in Figure 2:

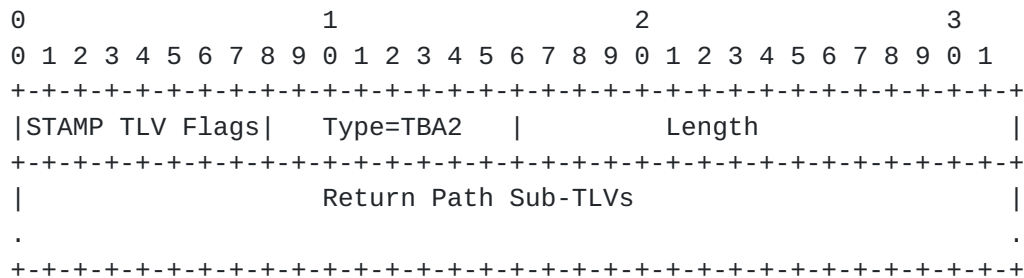


Figure 2: Return Path TLV

The STAMP TLV Flags are set using the procedures described in [\[RFC8972\]](#).

The Return Path TLV is optional. The STAMP Session-Sender MUST only insert one Return Path TLV in the STAMP test packet. The STAMP Session-Reflector that supports this TLV, MUST only process the first Return Path TLV in the test packet and ignore other Return Path TLVs if present, and it MUST NOT add Return Path TLV in the reply test packet.

4.1. Return Path Sub-TLVs

The Return Path TLV contains one or more Sub-TLVs to carry the information for the requested return path. A Return Path Sub-TLV can either carry Return Path Control Code, Return Path IP Address or Return Path Segment List.

The STAMP Sub-TLV Flags are set using the procedures described in [\[RFC8972\]](#).

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

4.1.1. 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:

- o Type (value 1): Return Path Control Code. The STAMP Session-Sender can request the STAMP 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: In-band Reply Requested.

When Control Code flag is set to 0x0 in the STAMP Session-Sender test packet, the Session-Reflector does not transmit reply test packet to the Session-Sender and terminates the STAMP test packet. 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 are ignored in this case.

When Control Code flag is set to 0x1 in the STAMP Session-Sender test packet, the Session-Reflector transmits the reply test packet in-band over the same incoming link where the test packet is received in the reverse direction.

[4.1.2.](#) Return Address Sub-TLV

The STAMP reply test packet may be transmitted to a different node than the Session-Sender (e.g. to a controller for telemetry use-cases). For this, the Session-Sender can specify in the test packet the receiving destination node address for the Session-Reflector reply test packet.

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:

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

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.

5. Security Considerations

The performance measurement is intended for deployment in well-managed private and service provider networks. As such, it assumes that a node involved in a measurement operation has previously verified the integrity of the path and the identity of the STAMP Session-Reflector.

If desired, attacks can be mitigated by performing basic validation and sanity checks, at the STAMP Session-Sender, of the timestamp fields in received measurement reply packets. The minimal state associated with these protocols also limits the extent of measurement disruption that can be caused by a corrupt or invalid packet to a single test cycle.

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

6. IANA Considerations

IANA will create a "STAMP TLV Type" registry for [\[RFC8972\]](#). IANA is requested to allocate a value for the following Destination Address TLV Type from the IETF Review TLV range of this registry. This TLV is to be carried in the STAMP test packets.

- o Type TBA1: Destination Node Address TLV

IANA is also requested to allocate a value for the following Return Path TLV Type from the IETF Review TLV range of the same registry. This TLV is to be carried in the STAMP test packets.

- o Type TBA2: Return Path TLV

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 1:

Value	Description	Reference
0	Reserved	This document
1 - 175	Unassigned	This document
176 - 239	Unassigned	This document
240 - 251	Experimental	This document
252 - 254	Private Use	This document
255	Reserved	This document

Table 1: Return Path Sub-TLV Type Registry

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

- o Type (value 1): Return Path Control Code
- o Type (value 2): Return Address
- o Type (value 3): SR-MPLS Label Stack of the Return Path
- o Type (value 4): SR-MPLS Binding SID of the Return SR Policy
- o Type (value 5): SRv6 Segment List of the Return Path
- o Type (value 6): SRv6 Binding SID of the Return SR Policy

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

[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>>.

[I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-09](#) (work in progress), November 2020.

[I-D.ietf-pce-binding-label-sid]
Sivabalan, S., Filsfils, C., Tantsura, J., Hardwick, J., Previdi, S., and C. Li, "Carrying Binding Label/Segment-ID in PCE-based Networks.", [draft-ietf-pce-binding-label-sid-05](#) (work in progress), October 2020.

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, 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

