

SPRING Working Group
Internet-Draft
Intended status: Informational
Expires: October 31, 2021

R. Gandhi, Ed.
C. Filsfils
Cisco Systems, Inc.
D. Voyer
Bell Canada
M. Chen
Huawei
B. Janssens
Colt
R. Foote
Nokia
April 29, 2021

**Performance Measurement Using Simple TWAMP (STAMP) for Segment Routing
Networks
draft-gandhi-spring-stamp-srpm-06**

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 describes procedures for Performance Measurement in SR networks using the mechanisms defined in [RFC 8762](#) (Simple Two-Way Active Measurement Protocol (STAMP)) and its optional extensions defined in [RFC 8972](#) and further augmented in [draft-gandhi-ippm-stamp-srpm](#). The procedure described is applicable to SR-MPLS and SRv6 data planes and is used for both links and end-to-end SR paths including SR Policies.

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 October 31, 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	3
2.	Conventions Used in This Document	3
2.1.	Abbreviations	3
2.2.	Reference Topology	4
3.	Overview	5
3.1.	Example STAMP Reference Model	6
4.	Delay Measurement for Links and SR Paths	7
4.1.	Session-Sender Test Packet	7
4.1.1.	Session-Sender Test Packet for Links	7
4.1.2.	Session-Sender Test Packet for SR Paths	8
4.2.	Session-Reflector Test Packet	10
4.2.1.	One-way Measurement Mode	11
4.2.2.	Two-way Measurement Mode	11
4.2.3.	Loopback Measurement Mode	13
4.3.	Delay Measurement for P2MP SR Policies	14
4.4.	Additional STAMP Test Packet Processing Rules	15
4.4.1.	TTL	15
4.4.2.	IPv6 Hop Limit	16
4.4.3.	Router Alert Option	16
4.4.4.	UDP Checksum	16
5.	Packet Loss Measurement for Links and SR Paths	16
6.	Direct Measurement for Links and SR Paths	16
7.	Session State for Links and SR Paths	17
8.	ECMP Support for SR Policies	17
9.	Security Considerations	18
10.	IANA Considerations	18
11.	References	19
11.1.	Normative References	19
11.2.	Informative References	19
	Acknowledgments	22
	Authors' Addresses	22

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 takes advantage of the Equal-Cost Multipaths (ECMPs) between source and transit nodes, between transit nodes and between transit and destination nodes. 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](#)] without the use of a control channel to pre-signal session parameters. [[RFC8972](#)] defines optional extensions for STAMP. [[I-D.gandhi-ippm-stamp-srpm](#)] augments that framework to define STAMP extensions for SR networks.

This document describes procedures for Performance Measurement in SR networks using the mechanisms defined in STAMP [[RFC8762](#)] and its optional extensions defined in [[RFC8972](#)] and further augmented in [[I-D.gandhi-ippm-stamp-srpm](#)]. The procedure described is applicable to SR-MPLS and SRv6 data planes and is used for both links and end-to-end SR paths including SR Policies [[RFC8402](#)].

2. Conventions Used in This Document

2.1. Abbreviations

BSID: Binding Segment ID.

DM: Delay Measurement.

ECMP: Equal Cost Multi-Path.

HMAC: Hashed Message Authentication Code.

LM: Loss Measurement.

MPLS: Multiprotocol Label Switching.

NTP: Network Time Protocol.

OWAMP: One-Way Active Measurement Protocol.

PM: Performance Measurement.

PSID: Path Segment Identifier.

PTP: Precision Time Protocol.

SHA: Secure Hash Algorithm.

SID: Segment ID.

SL: Segment List.

SR: Segment Routing.

SRH: Segment Routing Header.

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.

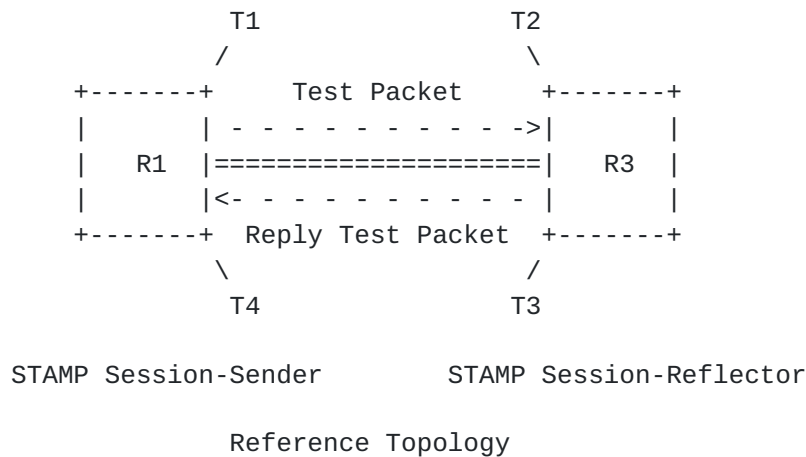
TC: Traffic Class.

TTL: Time To Live.

2.2. 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 may be transmitted to the STAMP Session-Sender R1 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.

The nodes R1 and R3 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 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. Overview

For performance measurement in SR networks, the STAMP Session-Sender and Session-Reflector test packets defined in [RFC8762] are used. They are used in one-way, two-way (i.e. round-trip) and loopback measurement modes. Note that one-way and round-trip are referred to in [RFC8762] and are further described in this document because of the introduction of loopback measurement mode in SR networks. The procedures defined in this document are also used to infer packet loss in SR networks.

The STAMP test packets are transmitted on the same path as the data traffic flow under measurement to measure the delay and packet loss experienced by the data traffic flow.

Typically, the STAMP test packets are transmitted along an IP path between a Session-Sender and a Session-Reflector to measure delay and packet loss along that IP path. Matching the forward and reverse direction paths for STAMP test packets, even for directly connected nodes is not guaranteed.

It may be desired in SR networks that the same path (same set of links and nodes) between the Session-Sender and Session-Reflector be used for the STAMP test packets in both directions. This is achieved by using the optional STAMP extensions for SR-MPLS and SRv6 networks specified in [I-D.gandhi-ippm-stamp-srpm]. The STAMP Session-Reflector uses the return path parameters for the reply test packet from the received STAMP test packet, as described in [I-D.gandhi-ippm-stamp-srpm]. This way signaling and maintaining dynamic SR network state for the STAMP sessions on the Session-Reflector are avoided.

The optional STAMP extensions defined in [RFC8972] are used for direct measurement packet loss in SR networks.

3.1. Example STAMP Reference Model

An example of a STAMP reference model with some of the typical measurement parameters including the Reflector UDP port for STAMP test session is shown in the following Figure 1:

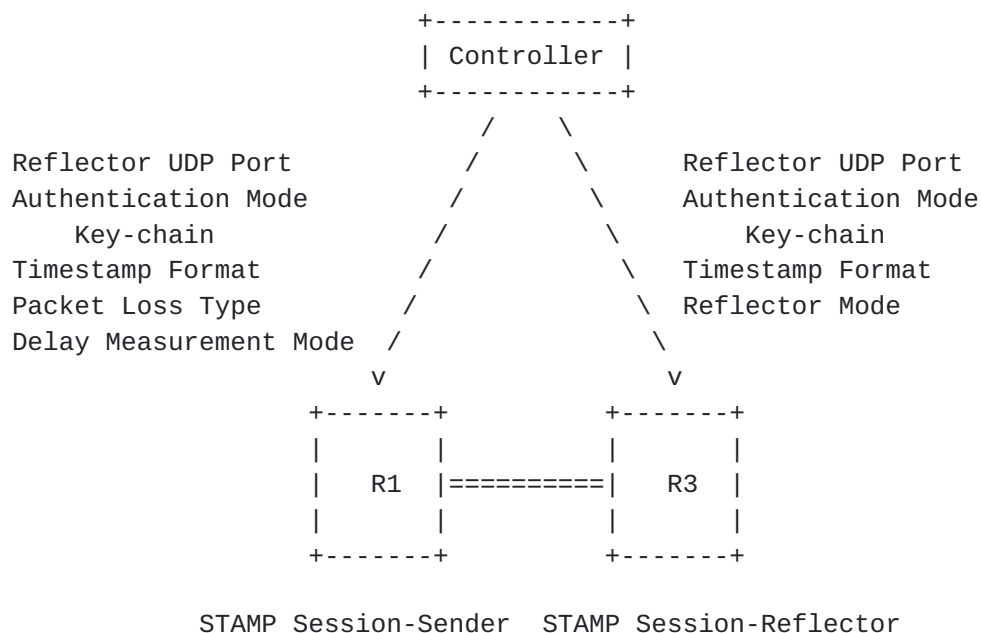


Figure 1: Example STAMP Reference Model

A reflector UDP port number is selected as described in [RFC8762]. The same reflector UDP port can be used for STAMP test sessions for link and end-to-end SR paths. In this case, the reflector UDP port does not distinguish between link or end-to-end SR path measurements.

Example of the Timestamp Format is Precision Time Protocol 64-bit truncated (PTPV2) [IEEE1588] and Network Time Protocol (NTP). By default, the Session-Reflector replies in kind to the timestamp format received in the received Session-Sender test packet, as indicated by the "Z" field in the Error Estimate field as described in [RFC8762].

The Session-Reflector mode can be Stateful or Stateless as defined in [RFC8762].

Example of Delay Measurement Mode is one-way, two-way (i.e. round-trip) and loopback mode as described in this document.

Example of Packet Loss Type can be round-trip, near-end (forward) and far-end (backward) packet loss as defined in [RFC8762].

When using the authenticated mode for the STAMP test sessions, the matching Authentication Type (e.g. HMAC-SHA-256) and Key-chain are user-configured on STAMP Session-Sender and STAMP Session-Reflector [RFC8762].

The controller shown in the example reference model is not intended for the dynamic signaling of the SR parameters for STAMP test sessions between the STAMP Session-Sender and STAMP Session-Reflector.

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.

4. Delay Measurement for Links and SR Paths

4.1. Session-Sender Test Packet

The content of an example STAMP Session-Sender test packet using an UDP header [RFC0768] is shown in Figure 2. The payload contains the STAMP Session-Sender test packet defined in [RFC8762].

```

+-----+
| IP Header |
. Source IP Address = Session-Sender IPv4 or IPv6 Address .
. Destination IP Address=Session-Reflector IPv4 or IPv6 Address.
. Protocol = UDP .
. .
+-----+
| UDP Header |
. Source Port = As chosen by Session-Sender .
. Destination Port = User-configured Reflector Port | 862 .
. .
+-----+
| Payload = Test Packet as specified in Section 4.2 of RFC 8762 |
. .
+-----+

```

Figure 2: Example Session-Sender Test Packet

4.1.1. Session-Sender Test Packet for Links

The STAMP Session-Sender test packet as shown in Figure 2 is transmitted over the link under delay measurement. The local and remote IP addresses of the link are used as Source and Destination Addresses, respectively. For IPv6 links, the link local addresses [RFC7404] can be used in the IPv6 header. The Session-Sender may use the local Address Resolution Protocol (ARP) table, Neighbor

Solicitation or other bootstrap method to find the IP address for the links and refresh. An IPv4 address from the range 127/8 or IPv6 loopback address ::1/128 [[RFC4291](#)] must not be used to IP route test packets in a network.

4.1.2. Session-Sender Test Packet for SR Paths

The delay measurement for end-to-end SR path in an SR network is applicable to both end-to-end SR-MPLS and SRv6 paths including SR Policies.

The STAMP Session-Sender IPv4 or IPv6 address is used as the Source Address. The SR Policy endpoint IPv4 or IPv6 address is used as the Destination Address.

In the case of Color-Only Destination Steering, with IPv4 endpoint of 0.0.0.0 or IPv6 endpoint of ::0 [[I-D.ietf-spring-segment-routing-policy](#)], the loopback address from the range 127/8 for IPv4, or the loopback address ::1/128 for IPv6 [[RFC4291](#)] is used as the Session-Reflector Address, respectively.

4.1.2.1. Session-Sender Test Packet for SR-MPLS Policies

An SR-MPLS Policy may contain a number of Segment Lists (SLs). A STAMP Session-Sender test packet is transmitted for each Segment List of the SR-MPLS Policy. The content of an example STAMP Session-Sender test packet for an end-to-end SR-MPLS Policy is shown in Figure 3.

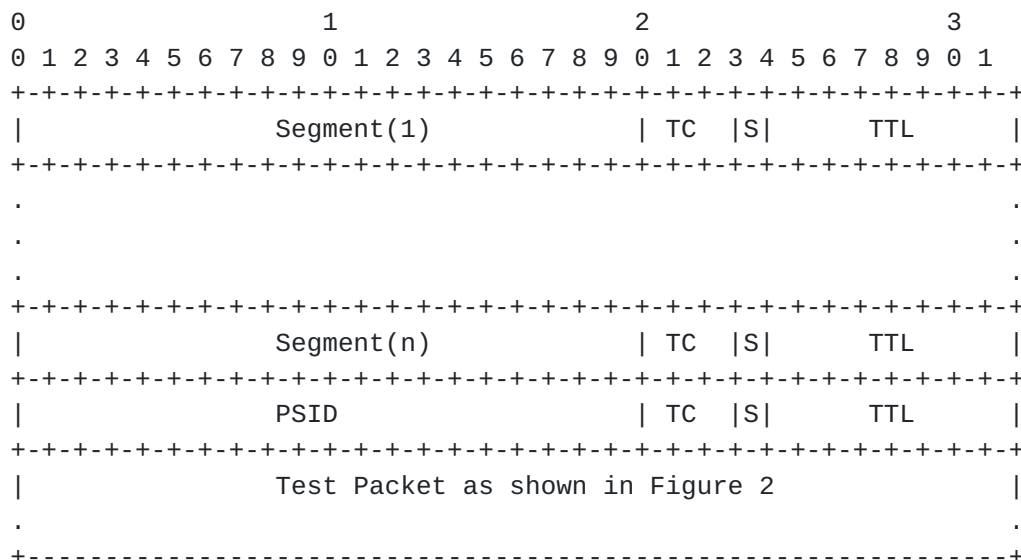


Figure 3: Example Session-Sender Test Packet for SR-MPLS Policy

The Segment List can be empty in case of a single-hop SR-MPLS Policy with Implicit NULL label.

The Path Segment Identifier (PSID)

[[I-D.ietf-spring-mpls-path-segment](#)] of an SR-MPLS Policy can be carried in the MPLS header as shown in Figure 3, and can be used for direct measurement as described in [Section 6](#), titled "Direct Measurement for Links and SR Paths".

[4.1.2.2](#). Session-Sender Test Packet for SRv6 Policies

An SRv6 Policy may contain a number of Segment Lists. A STAMP Session-Sender test packet is transmitted for each Segment List of the SRv6 Policy. An SRv6 Policy can contain an SRv6 Segment Routing Header (SRH) carrying a Segment List as described in [[RFC8754](#)]. The content of an example STAMP Session-Sender test packet for an end-to-end SRv6 Policy is shown in Figure 4.

The SRv6 network programming is described in [[RFC8986](#)]. The procedure defined for Upper-Layer Header processing for SRv6 End SIDs in [Section 4.1.1 in \[RFC8986\]](#) is used to process the IPv6/UDP header in the received test packets on the Session-Reflector.


```

+-----+
| IP Header |
. Source IP Address = Session-Sender IPv6 Address .
. Destination IP Address = Destination IPv6 Address .
. . .
+-----+
| SRH as specified in RFC 8754 |
. <PSID, Segment List> .
. . .
+-----+
| IP Header |
. Source IP Address = Session-Sender IPv6 Address .
. Destination IP Address = Session-Reflector IPv6 Address .
. Protocol = UDP .
. . .
+-----+
| UDP Header |
. Source Port = As chosen by Session-Sender .
. Destination Port = User-configured Reflector Port | 862 .
. . .
+-----+
| Payload = Test Packet as specified in Section 4.2 of RFC 8762 |
. . .
+-----+

```

Figure 4: Example Session-Sender Test Packet for SRv6 Policy

The Segment List (SL) may be empty and no SRH may be carried.

The Path Segment Identifier (PSID)

[[I-D.ietf-spring-srv6-path-segment](#)] of the SRv6 Policy can be carried in the SRH as shown in Figure 4 and can be used for direct measurement as described in [Section 6](#), titled "Direct Measurement for Links and SR Paths".

[4.2. Session-Reflector Test Packet](#)

The STAMP Session-Reflector reply test packet uses the IP/UDP information from the received test packet as shown in Figure 5.


```

+-----+
| IP Header |
. Source IP Address = Session-Reflector IPv4 or IPv6 Address .
. Destination IP Address .
.           = Source IP Address from Received Test Packet .
. Protocol = UDP .
. .
+-----+
| UDP Header |
. Source Port = As chosen by Session-Reflector .
. Destination Port = Source Port from Received Test Packet .
. .
+-----+
| Payload = Test Packet as specified in Section 4.3 of RFC 8762 |
. .
+-----+

```

Figure 5: Example Session-Reflector Test Packet

[4.2.1.](#) One-way Measurement Mode

In one-way delay measurement mode, a reply test packet as shown in Figure 5 is transmitted by the STAMP Session-Reflector, for both links and end-to-end SR Policies. The reply test packet may be transmitted on the same path or a different path in the reverse direction.

The STAMP Session-Sender address may not be reachable via IP route from the STAMP Session-Reflector. The STAMP Session-Sender in this case can send its reachability path information to the STAMP Session-Reflector using the Return Path TLV defined in [\[I-D.gandhi-ippm-stamp-srpm\]](#).

In this mode, as per Reference Topology, all timestamps T1, T2, T3, and T4 are collected by the test packets. However, only timestamps T1 and T2 are used to measure one-way delay as $(T2 - T1)$. The one-way delay measurement mode requires the clock on the Session-Sender and Session-Reflector to be synchronized.

[4.2.2.](#) Two-way Measurement Mode

In two-way (i.e. round-trip) delay measurement mode, a reply test packet as shown in Figure 5 is transmitted by the STAMP Session-Reflector on the same path in the reverse direction, e.g. on the reverse direction link or associated reverse SR path [\[I-D.ietf-pce-sr-bidir-path\]](#).

For two-way delay measurement mode for links, the STAMP Session-Reflector needs to transmit the reply test packet on the same link where the test packet is received. The STAMP Session-Sender can request in the test packet to the STAMP Session-Reflector to transmit the reply test packet back on the same link using the Control Code Sub-TLV in the Return Path TLV defined in [\[I-D.gandhi-ippm-stamp-srpm\]](#).

For two-way delay measurement mode for end-to-end SR paths, the STAMP Session-Reflector needs to transmit the reply test packet on a specific reverse path. The STAMP Session-Sender can request in the test packet to the STAMP Session-Reflector to transmit the reply test packet back on a given reverse path using a Segment List sub-TLV in the Return Path TLV defined in [\[I-D.gandhi-ippm-stamp-srpm\]](#).

In this mode, as per Reference Topology, all timestamps T1, T2, T3, and T4 are collected by the test packets. All four timestamps are used to measure two-way delay as $((T4 - T1) - (T3 - T2))$. When clock synchronization on the Session-Sender and Session-Reflector nodes is not possible, the one-way delay can be derived using two-way delay divided by two.

[4.2.2.1](#). Session-Reflector Test Packet for SR-MPLS Policies

The content of an example STAMP Session-Reflector reply test packet transmitted on the same path as the data traffic flow under measurement for two-way delay measurement of an end-to-end SR-MPLS Policy is shown in Figure 6.

```

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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Segment(1)          | TC |S|          TTL          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
.
.
.
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Segment(n)          | TC |S|          TTL          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Test Packet as shown in Figure 5          |
.
+-----+

```

Figure 6: Example Session-Reflector Test Packet for SR-MPLS Policy

4.2.2.2. Session-Reflector Test Packet for SRv6 Policies

The content of an example STAMP Session-Reflector reply test packet transmitted on the same path as the data traffic flow under measurement for two-way delay measurement of an end-to-end SRv6 Policy with SRH is shown in Figure 7.

The procedure defined for Upper-Layer Header processing for SRv6 End SIDs in [Section 4.1.1 in \[RFC8986\]](#) is used to process the IPv6/UDP header in the received reply test packets on the Session-Sender.

```
+-----+
| IP Header |
. Source IP Address = Session-Reflector IPv6 Address .
. Destination IP Address = Destination IPv6 Address .
. . .
+-----+
| SRH as specified in RFC 8754 |
. <Segment List> .
. . .
+-----+
| IP Header |
. Source IP Address = Session-Reflector IPv6 Address .
. Destination IP Address .
. = Source IPv6 Address from Received Test Packet .
. Protocol = UDP .
. . .
+-----+
| UDP Header |
. Source Port = As chosen by Session-Reflector .
. Destination Port = Source Port from Received Test Packet .
. . .
+-----+
| Payload = Test Packet as specified in Section 4.3 of RFC 8762 |
. . .
+-----+
```

Figure 7: Example Session-Reflector Test Packet for SRv6 Policy

4.2.3. Loopback Measurement Mode

The STAMP Session-Sender test packets are transmitted in loopback mode to measure loopback delay of a bidirectional circular path. In this mode, the received Session-Sender test packets are not punted out of the fast path in forwarding (to slow path or control-plane) at the STAMP Session-Reflector. In other words, the Session-Reflector does not process them and generate reply test packets.

The IP header of the STAMP Session-Sender test packet contains the Destination Address equals to the STAMP Session-Sender address and the Source Address equals to the STAMP Session-Reflector address. The Session-Sender sets the Reflector UDP port that it uses to receive the test packet. Optionally, the STAMP Session-Sender test packet can carry the return path information (e.g. return path label stack for SR-MPLS) as part of the SR header.

The Session-Sender can use the SSID field in the reply test packet and/ or local configuration to know that the test session is using the loopback mode. As the reply test packet is not generated by the STAMP Session-Reflector, the STAMP Session-Sender ignores the 'Session-Sender Sequence Number', 'Session-Sender Timestamp', 'Session-Sender Error Estimate', and 'Session-Sender TTL' in the received test packet. The Session-Sender sets these fields to 0 upon transmission.

In this mode, as per Reference Topology, the timestamps T1 and T4 are collected by the test packets. Both these timestamps are used to measure loopback delay as $(T4 - T1)$. When STAMP capability on the Session-Reflector node is not possible, the one-way delay can be derived using loopback delay divided by two. In this mode, the responder node processing time component reflects only the time required to loop the test packet from the incoming interface to the outgoing interface in forwarding plane.

4.3. Delay Measurement for P2MP SR Policies

The Point-to-Multipoint (P2MP) SR path that originates from a root node terminates on multiple destinations called leaf nodes (e.g. P2MP SR Policy [[I-D.ietf-pim-sr-p2mp-policy](#)]).

The procedures for delay and loss measurement described in this document for end-to-end P2P SR Policies are also equally applicable to the P2MP SR Policies. The procedure for one-way measurement is defined as following:

- o The STAMP Session-Sender root node transmits test packets using the Tree-SID defined in [[I-D.ietf-pim-sr-p2mp-policy](#)] for the P2MP SR-MPLS Policy as shown in Figure 8. The STAMP Session-Sender test packets may contain the replication SID as defined in [[I-D.ietf-spring-sr-replication-segment](#)].
- o The Destination Address is set to the loopback address from the range 127/8 for IPv4, or the loopback address ::1/128 for IPv6.
- o Each STAMP Session-Reflector leaf node transmits its node address in the Source Address of the reply test packets shown in Figure 5.

This allows the STAMP Session-Sender root node to identify the STAMP Session-Reflector leaf nodes of the P2MP SR Policy.

- o The P2MP root node measures the delay for each P2MP leaf node individually.

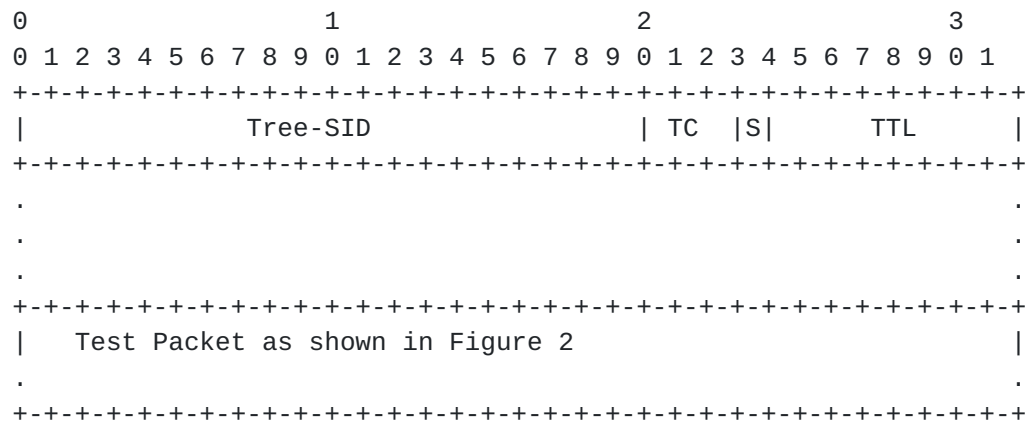


Figure 8: Example Session-Sender Test Packet with Tree-SID for SR-MPLS Policy

The considerations for two-way mode for P2MP SR Policy (e.g. for co-routed bidirectional SR-MPLS path) are outside the scope of this document.

[4.4.](#) Additional STAMP Test Packet Processing Rules

The processing rules described in this section are applicable to the STAMP test packets for links and end-to-end SR paths including SR Policies.

[4.4.1.](#) TTL

The TTL field in the IPv4 and MPLS headers of the STAMP Session-Sender and STAMP Session-Reflector test packets is set to 255, except in the following cases.

When using the Session-Reflector IPv4 Address from the range 127/8, the TTL field in the IPv4 header is set to 1, for otherwise, encapsulated packets.

For link delay, the TTL field in the STAMP test packet is set to 1 in one-way and two-way delay measurement modes.

4.4.2. IPv6 Hop Limit

The Hop Limit field in the IPv6 and SRH headers of the STAMP Session-Sender and STAMP Session-Reflector test packets is set to 255, except in the following cases.

When using the Session-Reflector IPv6 Address of loopback address `::1/128`, the Hop Limit field in the IPv6 header is set to 1, for otherwise, encapsulated packets.

For link delay, the Hop Limit field in the STAMP test packet is set to 1 in one-way and two-way delay measurement modes.

4.4.3. Router Alert Option

The Router Alert IP option (RAO) [[RFC2113](#)] is not set in the STAMP test packets for links and end-to-end SR paths.

4.4.4. UDP Checksum

For IPv4 test packets, where the hardware is not capable of re-computing the UDP checksum or adding checksum complement [[RFC7820](#)], the Session-Sender may set the UDP checksum value to 0 [[RFC8085](#)].

For IPv6 test packets, where the hardware is not capable of re-computing the UDP checksum or adding checksum complement [[RFC7820](#)], the Session-Sender and Session-Reflector may use the procedure defined in [[RFC6936](#)] for the UDP checksum.

5. Packet Loss Measurement for Links and SR Paths

The procedure described in [Section 4](#) for delay measurement using STAMP test packets can be used to detect (test) packet loss for links and end-to-end SR paths. The Sequence Number field in the STAMP test packet is used as described in [Section 4](#) "Theory of Operation" where Stateful and Stateless Session-Reflector operations are defined [[RFC8762](#)], to detect round-trip, near-end (forward) and far-end (backward) packet loss.

This method can be used for inferred packet loss measurement, however, it does not provide accurate data packet loss metric.

6. Direct Measurement for Links and SR Paths

The STAMP "Direct Measurement" TLV (Type 5) defined in [[RFC8972](#)] can be used in SR networks for data packet loss measurement. The STAMP test packets with this TLV are transmitted using the procedures

described in [Section 4](#) to collect the transmit and receive counters of the data flow for the links and end-to-end SR paths.

The PSID carried in the received data packet for the traffic flow under measurement can be used to measure receive data packets (for receive traffic counter) for an end-to-end SR path on the STAMP Session-Reflector. The PSID in the received Session-Sender test packet header can be used to associate the receive traffic counter on the Session-Reflector for the end-to-end SR path.

The STAMP "Direct Measurement" TLV (Type 5) lacks the support to identify the Block Number of the Direct Measurement traffic counters, which is required for Alternate-Marking Method [[RFC8321](#)] for accurate data packet loss metric.

7. Session State for Links and SR Paths

The STAMP test session state allows to know if the performance measurement test is active. The threshold-based notification may not be generated if the delay values do not change significantly. For an unambiguous monitoring, the controller needs to distinguish the cases whether the performance measurement is active, or delay values are not changing to cross threshold.

The STAMP test session state initially is declared active when one or more reply test packets are received at the STAMP Session-Sender. The STAMP test session state is declared idle (or failed) when consecutive N number of reply test packets are not received at the STAMP Session-Sender, where N is locally provisioned value.

8. ECMP Support for SR Policies

An SR Policy can have ECMPs between the source and transit nodes, between transit nodes and between transit and destination nodes. Usage of Anycast SID [[RFC8402](#)] by an SR Policy can result in ECMP paths via transit nodes part of that Anycast group. The test packets need to be transmitted to traverse different ECMP paths to measure end-to-end delay of an SR Policy.

Forwarding plane has various hashing functions available to forward packets on specific ECMP paths. The mechanisms described in [[RFC8029](#)] and [[RFC5884](#)] for handling ECMPs are also applicable to the delay measurement.

In IPv4 header of the STAMP Session-Sender test packets, sweeping of Session-Reflector Address from the range 127/8 can be used to exercise ECMP paths. In this case, both the forward and the return paths must be SR-MPLS paths when using the loopback mode.

As specified in [[RFC6437](#)], Flow Label field in the outer IPv6 header can also be used for sweeping to exercise different IPv6 ECMP paths.

The "Destination Node Address" TLV [[I-D.gandhi-ippm-stamp-srpm](#)] can be carried in the STAMP Session-Sender test packet to identify the intended Session-Reflector, for example, in case of using IPv4 Session-Reflector Address from 127/8 range when the STAMP test packet is encapsulated by a tunneling protocol or an MPLS Segment list. The STAMP Session-Reflector must not transmit reply test packet if it is not the intended destination node in the "Destination Node Address" TLV [[I-D.gandhi-ippm-stamp-srpm](#)].

9. 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 far-end STAMP Session-Reflector.

If desired, attacks can be mitigated by performing basic validation and sanity checks, at the STAMP Session-Sender, of the counter or timestamp fields in received measurement 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 packet to a single test cycle.

Use of HMAC-SHA-256 in the authenticated mode protects the data integrity of the test packets. SRv6 has HMAC protection authentication defined for SRH [[RFC8754](#)]. Hence, test packets for SRv6 may not need authentication mode. Cryptographic measures may be enhanced by the correct configuration of access-control lists and firewalls.

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

When using the procedures defined in [[RFC6936](#)], the security considerations specified in [[RFC6936](#)] also apply.

10. IANA Considerations

This document does not require any IANA action.

11. References

11.1. Normative References

- [RFC0768] Postel, J., "User Datagram Protocol", STD 6, [RFC 768](#), DOI 10.17487/RFC0768, August 1980, <<https://www.rfc-editor.org/info/rfc768>>.
- [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>>.
- [I-D.gandhi-ippm-stamp-srpm]
Gandhi, R., Filsfils, C., Voyer, D., Chen, M., and B. Janssens, "Simple TWAMP (STAMP) Extensions for Segment Routing Networks", [draft-gandhi-ippm-stamp-srpm-03](#) (work in progress), April 2021.

11.2. Informative References

- [IEEE1588]
IEEE, "1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", March 2008.
- [RFC2113] Katz, D., "IP Router Alert Option", [RFC 2113](#), DOI 10.17487/RFC2113, February 1997, <<https://www.rfc-editor.org/info/rfc2113>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", [RFC 5884](#), DOI 10.17487/RFC5884, June 2010, <<https://www.rfc-editor.org/info/rfc5884>>.

- [RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", [RFC 6437](#), DOI 10.17487/RFC6437, November 2011, <<https://www.rfc-editor.org/info/rfc6437>>.
- [RFC6936] Fairhurst, G. and M. Westerlund, "Applicability Statement for the Use of IPv6 UDP Datagrams with Zero Checksums", [RFC 6936](#), DOI 10.17487/RFC6936, April 2013, <<https://www.rfc-editor.org/info/rfc6936>>.
- [RFC7404] Behringer, M. and E. Vyncke, "Using Only Link-Local Addressing inside an IPv6 Network", [RFC 7404](#), DOI 10.17487/RFC7404, November 2014, <<https://www.rfc-editor.org/info/rfc7404>>.
- [RFC7820] Mizrahi, T., "UDP Checksum Complement in the One-Way Active Measurement Protocol (OWAMP) and Two-Way Active Measurement Protocol (TWAMP)", [RFC 7820](#), DOI 10.17487/RFC7820, March 2016, <<https://www.rfc-editor.org/info/rfc7820>>.
- [RFC8029] Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", [RFC 8029](#), DOI 10.17487/RFC8029, March 2017, <<https://www.rfc-editor.org/info/rfc8029>>.
- [RFC8321] Fioccola, G., Ed., Capello, A., Cociglio, M., Castaldelli, L., Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi, "Alternate-Marking Method for Passive and Hybrid Performance Monitoring", [RFC 8321](#), DOI 10.17487/RFC8321, January 2018, <<https://www.rfc-editor.org/info/rfc8321>>.
- [RFC8085] Eggert, L., Fairhurst, G., and G. Shepherd, "UDP Usage Guidelines", [BCP 145](#), [RFC 8085](#), DOI 10.17487/RFC8085, March 2017, <<https://www.rfc-editor.org/info/rfc8085>>.
- [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>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", [RFC 8754](#), DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.

- [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>>.
- [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-spring-sr-replication-segment] Voyer, D., Filsfils, C., Parekh, R., Bidgoli, H., and Z. Zhang, "SR Replication Segment for Multi-point Service Delivery", [draft-ietf-spring-sr-replication-segment-04](#) (work in progress), February 2021.
- [I-D.ietf-pim-sr-p2mp-policy] Voyer, D., Filsfils, C., Parekh, R., Bidgoli, H., and Z. Zhang, "Segment Routing Point-to-Multipoint Policy", [draft-ietf-pim-sr-p2mp-policy-02](#) (work in progress), February 2021.
- [I-D.ietf-spring-mpls-path-segment] Cheng, W., Li, H., Chen, M., Gandhi, R., and R. Zigler, "Path Segment in MPLS Based Segment Routing Network", [draft-ietf-spring-mpls-path-segment-04](#) (work in progress), April 2021.
- [I-D.ietf-spring-srv6-path-segment] Li, C., Cheng, W., Chen, M., Dhody, D., and R. Gandhi, "Path Segment for SRv6 (Segment Routing in IPv6)", [draft-ietf-spring-srv6-path-segment-00](#) (work in progress), November 2020.
- [I-D.ietf-pce-sr-bidir-path] Li, C., Chen, M., Cheng, W., Gandhi, R., and Q. Xiong, "Path Computation Element Communication Protocol (PCEP) Extensions for Associated Bidirectional Segment Routing (SR) Paths", [draft-ietf-pce-sr-bidir-path-05](#) (work in progress), January 2021.
- [I-D.ietf-ippm-stamp-yang] Mirsky, G., Min, X., and W. Luo, "Simple Two-way Active Measurement Protocol (STAMP) Data Model", [draft-ietf-ippm-stamp-yang-07](#) (work in progress), March 2021.

[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, Gyan Mishra, Xie Jingrong, and Mike Koldychev for reviewing this document and providing useful comments and suggestions. Patrick Khordoc and Radu Valceanu have helped improve the mechanisms described in this document.

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