IPPM Working Group R. Gandhi, Ed.
Internet-Draft C. Filsfils
Intended status: Standards Track Cisco Systems, Inc.
Expires: 5 February 2024 M. Chen
Huawei
B. Janssens
Colt
R. Foote
Nokia
4 August 2023

Simple TWAMP (STAMP) Extensions for Segment Routing Networks
draft-ietf-ippm-stamp-srpm-18

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

Copyright Notice

Table of Contents

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", "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

MPLS: Multiprotocol Label Switching.

SID: Segment Identifier.

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 T1 is a transmit timestamp and T4 is a receive timestamp added by node S1 in the STAMP test packet. The T2 is a receive timestamp and T3 is a transmit timestamp added by node R1 in the STAMP test packet.

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).

```
                       T1                T2
                      /                   \
             +-------+     Test Packet     +-------+
             |       | - - - - - - - - - ->|       |
             |   S1  |=====================|   R1  |
             |       |<- - - - - - - - - - |       |
             +-------+  Reply Test Packet  +-------+
                      \                   /
                       T4                T3

         STAMP Session-Sender        STAMP Session-Reflector
```

Reference Topology

3. Destination Node Address TLV

The Session-Sender may need to transmit test packets to the Session-Reflector with a destination address that is not a routable (i.e., suitable for use as the Source Address of the reply test packet) address of the Session-Reflector. This can be facilitated, for example, by encapsulating the STAMP packet by a tunneling protocol, see Appendix A, for a worked example.

[RFC8972] defines STAMP Session-Sender and Session-Reflector test packets that can include one or more optional TLVs. In this document, the TLV type (value 9 for IPv4 and IPv6) is defined for the Destination Node Address TLV for the STAMP test packet [RFC8972]. The formats of the Destination Node Address TLVs are shown in Figure 1:

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|STAMP TLV Flags|    Type=9     |         Length=4              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          IPv4 Address                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|STAMP TLV Flags|    Type=9     |         Length=16             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                          IPv6 Address                         |
|                                                               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 1: Destination Node Address TLV Format

TLV fields are defined as follows:

STAMP TLV Flags : The STAMP TLV Flags follow the procedures described in [RFC8972] and this document.

Type : Type (value 9) for IPv4 Destination Node Address TLV or IPv6 Destination Node Address TLV.

Length : A two-octet field equal to the length of the Address field in octets. The length is 4 octets for IPv4 address and 16 octets for IPv6 address.

The Destination Node Address TLV indicates an address of the intended Session-Reflector node of the test packet. If the received Destination Node Address is one of the addresses of the Session-Reflector, it SHOULD be used as the Source Address in the IP header of the reply test packet. If the Destination Node Address TLV is sent, the SSID MUST also be sent.

A Session-Reflector that recognizes this TLV, MUST set the U flag [RFC8972] 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. In this case, the Session-Reflector does not use the received Destination Node Address as the Source Address in the IP header of the reply test packet. Otherwise, the Session-Reflector MUST set the U flag in the Destination Node Address TLV in the reply test packet to 0.

4. 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, for measurement 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:

```
    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
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |STAMP TLV Flags|    Type=10    |         Length                |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                   Return Path Sub-TLVs                        |
    .                                                               .
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2: Return Path TLV

TLV fields are defined as follows:

STAMP TLV Flags : The STAMP TLV Flags follow the procedures described in [RFC8972] and this document.

Type : Type (value 10) for Return Path TLV.

Length : A two-octet field equal to the length of the Return Path Sub-TLVs field in octets.

Return Path Sub-TLVs : As defined in Section 4.1.

A Session-Sender MUST NOT insert more than one Return Path TLV in the STAMP test packet. A 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. A Session-Reflector that supports this TLV MUST reply using the Return Path received in the Session-Sender test packet, if no error was encountered while processing the TLV.

A Session-Reflector that recognizes this TLV, MUST set the U flag [RFC8972] in the reply test packet to 1 if the Session-Reflector determined that it cannot use the return path in the test packet to transmit the reply test packet. Otherwise, the Session-Reflector MUST set the U flag in the reply test packet to 0.

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 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].

A Return Path TLV MUST NOT contain more than one Control Code Sub-TLV or more than one Return Address Sub-TLV or more than one Segment List Sub-TLV in Session-Sender test packet.

A Return Path TLV MUST NOT contain both Control Code Sub-TLV as well as Return Address or Return Segment List Sub-TLV in Session-Sender test packet.

A Return Path TLV MAY contain both Return Address as well as Return Segment List Sub-TLV in Session-Sender test packet.

4.1.1. Return Path Control Code Sub-TLV

The format of the Return Path Control Code Sub-TLV is shown in Figure 3.

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|STAMP TLV Flags|    Type=1     |          Length=4             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                   Control Code Flags                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

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

TLV fields are defined as follows:

* 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 Flags field.

STAMP TLV Flags : The STAMP TLV Flags follow the procedures described in [RFC8972] and this document.

Length : A two-octet field equal to the length of the Control Code flags which is 4 octets.

Control Code Flags (32-bit): Reply Request Flag at bit 31 (least significant bit) is defined as follows.

    0x0 : No Reply Requested.

    0x1 : Reply Requested on the Same Link.

All other bits are reserved and must be transmitted as 0 and ignored by the receiver.

When Control Code flag for Reply Request 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 for Reply Request 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. The link may be a physical interface, virtual link, or Link Aggregation Group (LAG) [IEEE802.1AX], or LAG member. All other Return Path Sub-TLVs MUST be ignored in this case. When using LAG member links, STAMP extension for Micro-Session ID TLV defined in [I-D.ietf-ippm-stamp-on-lag] can be used to identify the link.

4.1.2. Return Address Sub-TLV

The STAMP reply test packet may be transmitted to the Session-Sender to the specified Return Address in the Return Address Sub-TLV instead of transmitting to the Source Address in the Session-Sender test packet.

The formats of the IPv4 and IPv6 Return Address Sub-TLVs are shown in Figure 4.

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|STAMP TLV Flags|     Type=2    |         Length=4              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Return IPv4 Address                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|STAMP TLV Flags|     Type=2    |         Length=16             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                    Return IPv6 Address                        |
|                                                               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

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

The TLV fields are defined as follows:

* Type : Type (value 2) for IPv4 Return Address or IPv6 Return Address.

The Return Address requests that the Session-Reflector reply test packet be sent to the specified address, rather than to the Source Address in the Session-Sender test packet.

STAMP TLV Flags : The STAMP TLV Flags follow the procedures described in [RFC8972] and this document.

Length : A two-octet field equal to the length of the Return Address field in octets. The length is 4 octets for IPv4 address and 16 octets for IPv6 address.

4.1.3. Return Segment List Sub-TLVs

The format of the Segment List Sub-TLVs in the Return Path TLV is shown in Figures 5 and 6. The Segments carried in Segment List Sub-TLVs are described in [RFC8402]. The segment entries MUST be in network order.

The Session-Sender MUST only insert one Segment List Return Path Sub-TLV in the test packet and Segment List MUST contain at least one Segment. 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.

TLV fields are defined as follows:

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

STAMP TLV Flags : The STAMP TLV Flags follow the procedures described in [RFC8972] and this document.

Length : A two-octet field equal to the length of the Segment List field in octets. Length MUST NOT be 0.

4.1.3.1. Return Path SR-MPLS Segment-List Sub-TLV

```
    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
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |STAMP TLV Flags|     Type=3    |         Length                |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |      Segment(1)                       | TC  |S|      TTL      |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    .                                                               .
    .                                                               .
    .                                                               .
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |      Segment(n) (bottom of stack)     | TC  |S|      TTL      |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5: SR-MPLS 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 Time-To-Live (TTL) value, 3-bit Traffic Class (TC) value and 1-bit End-Of-Stack (S) field. Length of the Sub-TLV modulo 4 MUST be 0.

As an example, an SR-MPLS Label Stack Sub-TLV could carry only the Binding SID Label [I-D.ietf-pce-binding-label-sid] of the Return SR-MPLS Policy. The Binding SID Label of the Return SR-MPLS Policy is local to the Session-Reflector. The mechanism to signal the Binding SID Label to the Session-Sender is outside the scope of this document.

As another example, an SR-MPLS Label Stack Sub-TLV could include the Path Segment Identifier Label of the Return SR-MPLS Policy in the Segment List of the SR-MPLS Policy.

4.1.3.2. Return Path SRv6 Segment-List Sub-TLV

```
     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
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |STAMP TLV Flags|     Type=4    |         Length                |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    |      Segment(1) (128-bit IPv6 address)                        |
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    .                                                               .
    .                                                               .
    .                                                               .
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    |      Segment(n) (128-bit IPv6 address) (bottom of stack)      |
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

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

The SRv6 Segment List contains a list of 128-bit IPv6 addresses representing the SRv6 SIDs. Length of the Sub-TLV modulo 16 MUST be 0.

As an example, an SRv6 Segment List Sub-TLV could carry only the SRv6 Binding SID [I-D.ietf-pce-binding-label-sid] of the Return SRv6 Policy. The SRv6 Binding SID of the Return SRv6 Policy is local to the Session-Reflector. The mechanism to signal the SRv6 Binding SID to the Session-Sender is outside the scope of this document.

As another example, an SRv6 Segment List Sub-TLV could include the SRv6 Path Segment Identifier of the Return SRv6 Policy in the Segment List of the SRv6 Policy.

5. Interoperability with TWAMP Light

This document does not introduce any additional considerations for interoperability with TWAMP Light than those described in Section 4.6 of [RFC8762].

As described in [RFC8762], there are two possible combinations for such an interoperability use case:

- STAMP Session-Sender with TWAMP Light Session-Reflector

- TWAMP Light Session-Sender with STAMP Session-Reflector

If any of STAMP extensions defined in this document are used by STAMP Session-Sender, the TWAMP Light Session-Reflector will view them as the Packet Padding field.

6. Security Considerations

The security considerations specified in [RFC8762] and [RFC8972] also apply to the extensions defined in this document. Specifically, the authenticated mode and 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 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.

If desired, attacks can be mitigated by performing basic validation checks of the timestamp fields (such as T2 is later than T1 in the Reference Topology in Section 2.3) in received reply test packets at the Session-Sender. The minimal state associated with these protocols also limit the extent of measurement disruption that can be caused by a corrupt or invalid test packet to a single test cycle.

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

The STAMP extensions defined in this document may be used for potential address spoofing. For example, a Session-Sender may specify a Return Path IP Address that is different from the Session-Sender address. The Session-Reflector MAY drop the Session-Sender test packet when it cannot determine whether the Return Path IP Address is local on the Session-Sender. To help Session-Reflector to make that determination, the Return Path IP Address may also be provisioned by the operator, for example, in an access control list.

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 IPv4 or IPv6 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
0 - 175	IETF Review	This document
176 - 239	First Come, First Served	This document
240 - 251	Experimental Use	This document
252 - 255	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 IPv4 or IPv6 Address	This document
3	SR-MPLS Label Stack of the Return Path	This document
4	SRv6 Segment List of the Return Path	This document
255	Reserved	This document

Table 3: Return Path Sub-TLV Types

IANA is requested to create a sub-registry for "Return Path Control Code Flags" for the Return Path Control Code Sub-TLV. All code points in the bit position 31 (counting from bit 31 as the least significant bit) through 12 in this registry shall be allocated according to the "IETF Review" procedure as specified in [RFC8126]. Code points in the bit position 11 through 8 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 4:

Bit	Description	Reference
31 - 12	IETF Review	This document
11 - 8	First Come, First Served	This document
7 - 4	Experimental Use	This document
3 - 0	Private Use	This document

Table 4: Return Path Control Code Flags Registry

IANA is requested to allocate the value for the following Return Path Control Code Flag from this registry.

Bit	Description	Reference
31	Reply Request	This document

Table 5: Return Path Control Code 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>.

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>.

[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-16, 27 March 2023, <https://www.ietf.org/archive/id/draft-ietf-pce-binding-label-sid-16.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-11, 13 March 2023, <https://www.ietf.org/archive/id/draft-ietf-ippm-stamp-yang-11.txt>.

[I-D.ietf-ippm-stamp-on-lag]
Li, Z., Zhou, T., Guo, J., Mirsky, G., and R. Gandhi, "Simple Two-Way Active Measurement Protocol Extensions for Performance Measurement on LAG", Work in Progress, Internet-Draft, draft-ietf-ippm-stamp-on-lag-03, 2 July 2023, <https://www.ietf.org/archive/id/draft-ietf-ippm-stamp-on-lag-03.txt>.

[IEEE802.1AX]
IEEE Std. 802.1AX, "IEEE Standard for Local and metropolitan area networks - Link Aggregation", November 2008.

Appendix A. Destination Node Address TLV Use-case Example

The STAMP test packets can be encapsulated with an SR-MPLS Segment List and IPv4 header containing destination IPv4 address from 127/8 range or STAMP test packets encapsulated with outer IPv6 header and Segment Routing Header (SRH) with inner IPv6 header containing IPv6 destination IPv6 address ::1/128.

In an ECMP environment, the hashing function in forwarding may decide the outgoing path using the source address, destination address, UDP ports, IPv6 flow-label, etc. from the packet. Hence, for IPv4, for example, different values of IPv4 destination address from 127/8 range may be used in the IPv4 header of the STAMP test packets to measure different ECMP paths. For IPv6, for example, different values of flow-label may be used in the IPv6 header of the STAMP test packets to measure different ECMP paths.

In those cases, the STAMP test packets may reach a node that is not the Session-Reflector for this STAMP session in an error condition, and this un-intended node may transmit reply test packet that can result in reporting of invalid measurement metrics. The intended Session-Reflector address can be carried in the Destination Node Address TLV to help detect this error.

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, Henrik Nydell, and Cheng Li for providing comments and suggestions. Thank you Joel Halpern for Gen-ART review, Martin Duke for AD review, and Kathleen Moriarty for Security review. The authors would like to thank Robert Wilton, Eric Vyncke, Paul Wouters, John Scudder, Roman Danyliw, and Jim Guichard for IESG review.

Contributors

The following people have substantially contributed to this document:

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

Authors' Addresses

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

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

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