SPRING Working Group
Internet-Draft
Intended Status: Standards Track
Expires: August 13, 2019

R. Gandhi, Ed.
C. Filsfils
Cisco Systems, Inc.
D. Voyer
Bell Canada
February 9, 2019

In-band Performance Measurement Using TWAMP for Segment Routing Networks
draft-gandhi-spring-twamp-srpm-00

Abstract

Segment Routing (SR) is applicable to both Multiprotocol Label Switching (SR-MPLS) and IPv6 (SRv6) data planes. This document specifies procedures for sending and processing in-band probe query and response messages for Performance Measurement. The procedure uses the mechanisms defined in RFC 5357 (Two-Way Active Measurement Protocol (TWAMP)) for Delay Measurement, and also uses the mechanisms for direct-mode loss measurement defined in this document. The procedure specified is applicable to SR-MPLS and SRv6 data planes for both links and end-to-end measurement for SR Policies.

Table of Contents

1. Introduction

Segment Routing (SR) technology greatly simplifies network operations for Software Defined Networks (SDNs). SR is applicable to both Multiprotocol Label Switching (SR-MPLS) and IPv6 (SRv6) data planes.

SR takes advantage of the Equal-Cost Multipaths (ECMPs) between source, transit and destination nodes. SR Policies as defined in [I-D.spring-segment-routing-policy] are used to steer traffic through a specific, user-defined path using a stack of Segments. Built-in SR Performance Measurement (PM) is one of the essential requirements to provide Service Level Agreements (SLAs).

The One-Way Active Measurement Protocol (OWAMP) defined in [RFC4656] and Two-Way Active Measurement Protocol (TWAMP) defined in [RFC5357] provide capabilities for the measurement of various performance metrics in IP networks. These protocols rely on control channel signaling to establish a test channel over an UDP path. These protocols lack support for direct-mode Loss Measurement (LM) to detect actual data traffic loss which is required in SR networks. The Simple Two-way Active Measurement Protocol (STAMP) [I-D.ippm-stamp] alleviates the control channel signaling by using configuration data model to provision test channels and required UDP ports. The TWAMP Light from broadband forum [BBF.TR-390] provides simplified mechanisms for active performance measurement in Customer Edge IP networks.

This document specifies procedures for sending and processing in-band probe query and response messages for Performance Measurement. The procedure uses the mechanisms defined in RFC 5357 (Two-Way Active Measurement Protocol (TWAMP)) for Delay Measurement, and also uses the mechanisms for direct-mode loss measurement defined in this document. The procedure specified is applicable to SR-MPLS and SRv6 data planes for both links and end-to-end measurement for SR Policies. For SR Policies, there are ECMPs between the source and transit nodes, between transit nodes and between transit and destination nodes. This document also defines mechanisms for handling ECMPs of SR Policies for performance delay measurement.

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

BSID: Binding Segment ID.

DM: Delay Measurement.

ECMP: Equal Cost Multi-Path.

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.

SID: Segment ID.

SL: Segment List.

SR: Segment Routing.

SR-MPLS: Segment Routing with MPLS data plane.

SRv6: Segment Routing with IPv6 data plane.

STAMP: Simple Two-way Active Measurement Protocol.

TC: Traffic Class.

TWAMP: Two-Way Active Measurement Protocol.

2.3. Reference Topology

In the reference topology, the querier node R1 initiates a probe query for performance measurement and the responder node R5 sends a probe response for the query message received. The probe response may be sent to the querier node R1. The nodes R1 and R5 may be directly connected via a link enabled with Segment Routing or there exists a Point-to-Point (P2P) SR Policy [I-D.spring-segment-routing-policy] on node R1 with destination to node R5. In case of Point-to-Multipoint (P2MP), SR Policy originating from source node R1 may terminate on multiple destination leaf nodes [I-D.spring-sr-p2mp-policy].

```
             +-------+        Query        +-------+
             |       | - - - - - - - - - ->|       |
             |   R1  |---------------------|   R5  |
             |       |<- - - - - - - - - - |       |
             +-------+       Response      +-------+
```

Reference Topology

Both Delay and Loss performance measurement is performed in-band for the traffic traversing between node R1 and node R5. One-way delay and two-way delay measurements are defined in [RFC4656] and [RFC5357], respectively. One-way loss measurement provides receive packet loss whereas two-way loss measurement provides both transmit and receive packet loss.

2.4. In-band Probe Messages

For both Delay and Loss measurements for links and SR Policies, no PM session is created on the responder node. The probe messages for Delay measurement are sent in-band by the querier node to measure the delay experienced by the actual traffic flowing on the links and SR Policies. For Loss measurement, in-band probe messages are used to collect the traffic counter for the incoming link or incoming SID on which the probe query message is received at the responder node R5 as it has no PM session state present on the node. The performance measurement for Delay and Loss using out-of-band probe query messages are outside the scope of this document.

3. Probe Messages

3.1. Probe Query Message

In this document, procedures using [RFC5357] is used for Delay and Loss measurements for SR links and end-to-end SR Policies. A user-configured UDP port is used for identifying PM probe packets that does not require to bootstrap PM sessions. A UDP port number from the Dynamic and/or Private Ports range 49152-65535 is used as the destination UDP port. This approach is similar to the one defined in STAMP protocol [I-D.ippm-stamp]. The IPv4 TTL or IPv6 Hop Limit field of the IP header MUST be set to 255.

3.1.1. Delay Measurement Probe Query Message

The message content for Delay Measurement probe query message using UDP header [RFC768] is shown in Figure 1. The DM probe query message is sent with user-configured Destination UDP port number [I-D.ippm-stamp]. The Source UDP port is set to the same value for two-way delay measurement. The DM probe query message contains the payload for delay measurement defined in Section 4.2.1 of [RFC5357] for TWAMP or in Section 4.1.2 of [RFC4656] for OWAMP.

```
+---------------------------------------------------------------+
| IP Header                                                     |
.  Source IP Address = Querier IPv4 or IPv6 Address             .
.  Destination IP Address = Responder IPv4 or IPv6 Address      .
.  Protocol = UDP                                               .
.  Router Alert Option Not Set                                  .
.                                                               .
+---------------------------------------------------------------+
| UDP Header                                                    |
.  Source Port = As chosen by Querier                           .
.  Destination Port = User-configured Port for Delay Measurement.
.                                                               .
+---------------------------------------------------------------+
| Payload = Message as specified in Section 4.2.1 of RFC 5357   |
| | Payload = Message as specified in Section 4.1.2 of RFC 4656 |
.                                                               .
+---------------------------------------------------------------+
```

Figure 1: DM Probe Query Message

Timestamp field is eight bytes and by default uses the IEEE 1588v2 Precision Time Protocol (PTP) truncated 64-bit timestamp format [IEEE1588].

3.1.2. Loss Measurement Probe Query Message

The message content for Loss Measurement probe query message using UDP header [RFC768] is shown in Figure 2. The LM probe query message is sent with user-configured Destination UDP port number [I-D.ippm-stamp]. The Source UDP port is set to the same value for two-way loss measurement. The LM probe query message contains the payload for loss measurement defined below.

```
+---------------------------------------------------------------+
| IP Header                                                     |
.  Source IP Address = Querier IPv4 or IPv6 Address             .
.  Destination IP Address = Responder IPv4 or IPv6 Address      .
.  Protocol = UDP                                               .
.  Router Alert Option Not Set                                  .
.                                                               .
+---------------------------------------------------------------+
| UDP Header                                                    |
.  Source Port = As chosen by Querier                           .
```

```
    .  Destination Port = User-configured Port for Loss Measurement .
    .                                                               .
    +---------------------------------------------------------------+
    |                        Sequence Number                        |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Transmit Counter                       |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Receive Counter                        |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Sender Sequence Number                 |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Sender Counter                         |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |  Sender TTL   |              Block Number                     |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Padding                                |
    .                                                               .
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2A: LM Probe Query Message for TWAMP

```
    +---------------------------------------------------------------+
    | IP Header                                                     |
    .  Source IP Address = Querier IPv4 or IPv6 Address             .
    .  Destination IP Address = Responder IPv4 or IPv6 Address      .
    .  Protocol = UDP                                               .
    .  Router Alert Option Not Set                                  .
    .                                                               .
    +---------------------------------------------------------------+
    | UDP Header                                                    |
    .  Source Port = As chosen by Querier                           .
    .  Destination Port = User-configured Port for Loss Measurement .
    .                                                               .
    +---------------------------------------------------------------+
    |                        Sequence Number                        |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        MBZ (12 octets)                        |
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Transmit Counter                       |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        MBZ (8 octets)                         |
```

```
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Receive Counter                        |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        MBZ (8 octets)                         |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Sender Sequence Number                 |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        MBZ (12 octets)                        |
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Sender Counter                         |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        MBZ (8 octets)                         |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |  Sender TTL   |              Block Number                     |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        MBZ (12 octets)                        |
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        HMAC (16 octets)                       |
    |                                                               |
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Padding                                |
    .                                                               .
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2B: LM Probe Query Message for TWAMP - Authenticated Mode

```
    +---------------------------------------------------------------+
    | IP Header                                                     |
    .  Source IP Address = Querier IPv4 or IPv6 Address             .
    .  Destination IP Address = Responder IPv4 or IPv6 Address      .
    .  Protocol = UDP                                               .
    .  Router Alert Option Not Set                                  .
    .                                                               .
    +---------------------------------------------------------------+
    | UDP Header                                                    |
    .  Source Port = As chosen by Querier                           .
```

```
    .  Destination Port = User-configured Port for Loss Measurement .
    .                                                               .
    +---------------------------------------------------------------+
    |                         Sequence Number                       |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                         Transmit Counter                      |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |  Sender TTL   |               Block Number                    |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                         Padding                               |
    .                                                               .
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2C: LM Probe Query Message for OWAMP

```
    +---------------------------------------------------------------+
    | IP Header                                                     |
    .  Source IP Address = Querier IPv4 or IPv6 Address             .
    .  Destination IP Address = Responder IPv4 or IPv6 Address      .
    .  Protocol = UDP                                               .
    .  Router Alert Option Not Set                                  .
    .                                                               .
    +---------------------------------------------------------------+
    | UDP Header                                                    |
    .  Source Port = As chosen by Querier                           .
    .  Destination Port = User-configured Port for Loss Measurement .
    .                                                               .
    +---------------------------------------------------------------+
    |                         Sequence Number                       |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                         MBZ (12 octets)                       |
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                         Transmit Counter                      |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                         MBZ (8 octets)                        |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |  Sender TTL   |               Block Number                    |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                         MBZ (12 octets)                       |
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

```
    |                        HMAC (16 octets)                       |
    |                                                               |
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                        Padding                                |
    .                                                               .
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2D: LM Probe Query Message for OWAMP - Authenticated Mode

Sequence Number (32-bit): As defined in [RFC5357].

Transmit Counter (64-bit): The number of packets sent by the querier node in the query message and by the responder node in the response message. The counter is always written at fixed location in the probe query and response messages.

Receive Counter (64-bit): The number of packets received at the responder node. It is written by the responder node in the probe response message.

Sender Counter (64-bit): This is the exact copy of the transmit counter from the received query message. It is written by the responder node in the probe response message.

Sender Sequence Number (32-bit): As defined in [RFC5357].

Sender TTL: As defined in [RFC5357].

Block Number (24-bit): The Loss Measurement using Alternate-Marking method defined in [RFC8321] requires to identify the Block Number (or color) of the traffic counters. The probe query and response messages carry Block Number for the traffic counters for loss measurement. In both probe query and response messages, the counters MUST belong to the same Block Number.

The Path Segment Identifier (PSID) [I-D.spring-mpls-path-segment] of the SR-MPLS Policy is used for accounting received traffic on the egress node for loss measurement.

3.1.3. Probe Query for SR Links

The probe query message as defined in Figure 1 is sent in-band for Delay measurement. The probe query message as defined in Figure 2 is sent in-band for Loss measurement.

3.1.4. Probe Query for End-to-end Measurement for SR Policy

3.1.4.1. Probe Query Message for SR-MPLS Policy

The message content for in-band probe query message using UDP header for end-to-end performance measurement of SR-MPLS Policy 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               Segment List(0)         | TC  |S|      TTL      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
.                                                               .
.                                                               .
.                                                               .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               Segment List(n)         | TC  |S|      TTL      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Message as shown in Figure 1 for DM or Figure 2 for LM      |
.                                                               .
+---------------------------------------------------------------+
```

Figure 3: Probe Query Message for SR-MPLS Policy

The Segment List (SL) can be empty to indicate Implicit NULL label case.

3.1.4.2. Probe Query Message for SRv6 Policy

The in-band probe query messages using UDP header for end-to-end performance measurement of an SRv6 Policy is sent using SRv6 Segment Routing Header (SRH) and Segment List of the SRv6 Policy as defined in [I-D.6man-segment-routing-header] and is 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            SRH                                |
.   END.OTP (DM) or END.OP (LM) with Target SRv6 SID            .
.                                                               .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Message as shown in Figure 1 for DM or Figure 2 for LM      |
.   (Using IPv6 Addresses)                                      .
.                                                               .
+---------------------------------------------------------------+
```

Figure 4: Probe Query Message for SRv6 Policy

For delay measurement of SRv6 Policy, END function END.OTP [I-D.spring-srv6-oam] is used with the target SRv6 SID to punt probe messages on the target node, as shown in Figure 4. Similarly, for loss measurement of SRv6 Policy, END function END.OP [I-D.spring-srv6-oam] is used with target SRv6 SID to punt probe messages on the target node.

3.2. Probe Response Message

The probe response message is sent using the IP/UDP information from the probe query message. The content of the probe response message is shown in Figure 5.

```
+---------------------------------------------------------------+
| IP Header                                                     |
.  Source IP Address = Responder IPv4 or IPv6 Address           .
.  Destination IP Address = Source IP Address from Query        .
.  Protocol = UDP                                               .
.  Router Alert Option Not Set                                  .
.                                                               .
+---------------------------------------------------------------+
| UDP Header                                                    |
.  Source Port = As chosen by Responder                         .
.  Destination Port = Source Port from Query                    .
.                                                               .
+---------------------------------------------------------------+
| Payload as specified in Section 4.2.1 of RFC 5357, or         |
. Payload as specified in Figure 2 in this document             .
.                                                               .
+---------------------------------------------------------------+
```

Figure 5: Probe Response Message

3.2.1. One-way Measurement

For one-way performance measurement, the probe response message as defined in Figure 5 is sent for both SR links and SR Policies.

3.2.2. Two-way Measurement

For two-way performance measurement, when using a bidirectional channel, the probe response message as defined in Figure 5 is sent back in-band to the querier node.

The Path Segment Identifier (PSID) [I-D.spring-mpls-path-segment] of the forward SR Policy can be used to find the reverse SR Policy to send the probe response message for two-way measurement of SR Policy.

3.2.2.1. Probe Response Message for SR-MPLS Policy

The message content for sending probe response message in-band using UDP header for two-way end-to-end performance measurement of an 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 List(0)        | TC  |S|      TTL      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
.                                                               .
.                                                               .
.                                                               .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                Segment List(n)        | TC  |S|      TTL      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                Message as shown in Figure 5                   |
.                                                               .
+---------------------------------------------------------------+
```

Figure 6: Probe Response Message for SR-MPLS Policy

3.2.2.2. Probe Response Message for SRv6 Policy

The message content for sending probe response message in-band using UDP header for two-way end-to-end performance measurement of an SRv6 Policy is shown in Figure 7.

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          SRH                                  |
.   END.OTP (DM) or END.OP (LM) with Target SRv6 SID            .
.                                                               .
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Message as shown in Figure 5 (with IPv6 Addresses)          |
.                                                               .
+---------------------------------------------------------------+
```

Figure 7: Probe Response Message for SRv6 Policy

4. Packet Loss Calculation

The formula for calculating the one-way packet loss using counters for a given block number is as following:

o One-way Packet_Loss[n-1, n] = (Sender_Counter[n] - Sender_Counter[n-1]) - (Receive_Counter[n] - Receive_Counter[n-1])

5. Performance Measurement for P2MP SR Policies

The procedures for delay and loss measurement described in this document for Point-to-Point (P2P) SR-MPLS Policies are also equally applicable to the Point-to-Multipoint (P2MP) SR Policies.

6. ECMP Support

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 PM probe messages need to be sent to traverse different ECMP paths to measure performance delay of an SR Policy.

Forwarding plane has various hashing functions available to forward packets on specific ECMP paths. Following mechanisms can be used in PM probe messages to take advantage of the hashing function in forwarding plane to influence the path taken by them.

o The mechanisms described in [RFC8029] [RFC5884] for handling ECMPs are also applicable to the performance measurement. In the IP/UDP header of the PM probe messages, Destination Addresses in 127/8 range for IPv4 or 0:0:0:0:0:FFFF:7F00/104 range for IPv6 can be used to exercise a particular ECMP path. As specified in [RFC6437], 3-tuple of Flow Label, Source Address and Destination Address fields in the IPv6 header can also be used.

o For SR-MPLS, entropy label [RFC6790] in the PM probe messages can be used.

o For SRv6, Flow Label in SRH [I-D.6man-segment-routing-header] of the PM probe messages can be used.

7. 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 responder node.

If desired, attacks can be mitigated by performing basic validation and sanity checks, at the querier, of the counter or timestamp fields in received measurement response messages. The minimal state associated with these protocols also limits the extent of measurement disruption that can be caused by a corrupt or invalid message to a single query/response cycle.

Use of HMAC-SHA-256 in the authenticated mode defined in this document protects the data integrity of the probe messages. SRv6 has HMAC protection authentication defined for SRH [I-D.6man-segment-routing-header]. Hence, PM probe messages for SRv6 may not need authentication mode. Cryptographic measures may be enhanced by the correct configuration of access-control lists and firewalls.

8. IANA Considerations

This document does not require any IANA actions.

9. References

9.1. Normative References

[RFC768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997.

[RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, September 2006.

[RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", RFC 8174, May 2017.

[I-D.spring-srv6-oam] Ali, Z., et al., "Operations, Administration, and Maintenance (OAM) in Segment Routing Networks with IPv6 Data plane (SRv6)", draft-ali-spring-srv6-oam.

9.2. Informative References

[IEEE1588] IEEE, "1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", March 2008.

[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.

[RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, November 2011.

[RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", RFC 6790, November 2012.

[RFC8029] Kompella, K., Swallow, G., Pignataro, C., Kumar, N., Aldrin, S. and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", RFC 8029, March 2017.

[RFC8321] Fioccola, G. Ed., "Alternate-Marking Method for Passive and Hybrid Performance Monitoring", RFC 8321, January 2018.

[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>.

[I-D.spring-segment-routing-policy] Filsfils, C., et al., "Segment Routing Policy Architecture", draft-ietf-spring-segment-routing-policy, work in progress.

[I-D.spring-sr-p2mp-policy] Voyer, D. Ed., et al., "SR Replication Policy for P2MP Service Delivery", draft-voyer-spring-sr-p2mp-policy, work in progress.

[I-D.spring-mpls-path-segment] Cheng, W., et al., "Path Segment in MPLS Based Segment Routing Network", draft-cheng-spring-mpls-path-segment, work in progress.

[I-D.6man-segment-routing-header] Filsfils, C., et al., "IPv6 Segment Routing Header (SRH)", draft-ietf-6man-segment-routing-header, work in progress.

[I-D.ippm-stamp] Mirsky, G. et al. "Simple Two-way Active

Measurement Protocol", draft-ietf-ippm-stamp, work in progress.

[BBF.TR-390] "Performance Measurement from IP Edge to Customer Equipment using TWAMP Light", BBF TR-390, May 2017.

Acknowledgments

TBA

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