

SPRING Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 13, 2020

R. Gandhi, Ed.
C. Filsfils
Cisco Systems, Inc.
D. Voyer
Bell Canada
S. Salsano
Universita di Roma "Tor Vergata"
M. Chen
Huawei
June 11, 2020

Performance Measurement Using [RFC 6374](#) with UDP Path for Segment Routing
Networks
[draft-gandhi-spring-rfc6374-srpm-udp-04](#)

Abstract

Segment Routing (SR) leverages the source routing paradigm. Segment Routing (SR) is applicable to both Multiprotocol Label Switching (SR-MPLS) and IPv6 (SRv6) data planes. This document specifies procedures for using UDP path for sending and processing probe query and response messages for Performance Measurement (PM). The procedure uses the mechanisms defined in [RFC 6374](#) for Performance Delay and Loss Measurement. The procedure specified is applicable to SR-MPLS and SRv6 data planes for both Links and end-to-end SR Paths including SR Policies measurements.

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 December 13, 2020.

Copyright Notice

Copyright (c) 2020 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	4
2.1.	Requirements Language	4
2.2.	Abbreviations	4
2.3.	Reference Topology	5
3.	Overview	5
3.1.	Example Provisioning Model	6
4.	Probe Query Message	6
4.1.	Delay Measurement Probe Query Message	6
4.2.	Loss Measurement Probe Query Message	7
4.3.	Combined Loss/Delay Measurement Probe Query Message	8
4.4.	Probe Query Message for Links	9
4.5.	Probe Query Message for SR Policies	9
4.5.1.	Probe Query Message for SR-MPLS Policy	9
4.5.2.	Probe Query Message for SRv6 Policy	10
5.	Probe Response Message	11
5.1.	One-way Measurement Mode	13
5.1.1.	Links and SR Policies	13
5.1.2.	Probe Response Message to Controller	13
5.2.	Two-way Measurement Mode	13
5.2.1.	Links	13
5.2.2.	SR Policies	13
5.2.3.	Return Path TLV	14
5.2.4.	Probe Response Message for SR-MPLS Policy	14
5.2.5.	Probe Response Message for SRv6 Policy	15
5.3.	Loopback Measurement Mode	15
6.	Performance Measurement for P2MP SR Policies	16
7.	ECMP Support for SR Policies	16
8.	Additional Probe Message Processing Rules	16
9.	Sequence Numbers	16
9.1.	Sequence Number TLV in Unauthenticated Mode	16

9.2.	Sequence Number TLV in Authenticated Mode	17
10.	Performance Delay and Liveness Monitoring	18
11.	Security Considerations	18
12.	IANA Considerations	19
13.	References	19
13.1.	Normative References	19
13.2.	Informative References	20
	Acknowledgments	22
	Contributors	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. 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).

[RFC6374] specifies protocol mechanisms to enable the efficient and accurate measurement of performance metrics and can be used in SR networks with MPLS data plane [[I-D.gandhi-mpls-rfc6374-sr](#)].

[RFC6374] addresses the limitations of the IP based performance measurement protocols as specified in [Section 1 of \[RFC6374\]](#).

[RFC6374] requires data plane to support MPLS Generic Associated Channel Label (GAL) and Generic Associated Channel (G-ACH), which may not be supported on all nodes in the segment routing network.

[RFC7876] specifies the procedures to be used when sending and processing out-of-band performance measurement probe response messages over an UDP return path for [RFC 6374](#) based probe queries.

[RFC7876] can be used to send out-of-band probe responses in both SR-MPLS and SRv6 networks for one-way performance measurement.

For SR Policies, there are ECMPs between the source and transit nodes, between transit nodes and between transit and destination nodes. [RFC 6374](#) does not define handling for ECMP forwarding paths when used in SR networks.

For two-way measurements for SR Policies, there is a requirement to specify a return path in the form of a Segment List in probe query messages that does not require on any SR Policy state information on the destination node.

This document specifies a procedure for sending and processing probe query and response messages using UDP paths for Performance Measurement in SR networks. The procedure uses [RFC 6374](#) defined mechanisms for Performance Delay and Loss Measurement and unless otherwise specified, the procedures from [RFC 6374](#) are not modified. The procedure specified is applicable to both SR-MPLS and SRv6 data planes. The procedure can be used for both Links and end-to-end SR Paths including SR Policies measurements.

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.

G-ACh: Generic Associated Channel (G-ACh).

GAL: Generic Associated Channel (G-ACh) Label.

LM: Loss Measurement.

MPLS: Multiprotocol Label Switching.

NTP: Network Time Protocol.

PM: Performance Measurement.

PSID: Path Segment Identifier.

PTP: Precision Time Protocol.

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.

TC: Traffic Class.

URO: UDP Return Object.

2.3. Reference Topology

In the reference topology shown below, the querier node R1 initiates a probe query for performance measurement and the responder node R5 sends a probe response for the probe query message received. The probe response may be sent to the querier node R1 or to a controller node R100. 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 Path e.g. SR Policy [[I-D.ietf-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.voyer-spring-sr-replication-segment](#)].



Reference Topology

3. Overview

For one-way, two-way and round-trip delay measurements, the procedures defined in [Section 2.4](#) and [Section 2.6 of \[RFC6374\]](#) are used. For transmit and receive packet loss measurements, the procedures defined in [Section 2.2](#) and [Section 2.6 of \[RFC6374\]](#) are used. The procedures use probe messages with UDP path and do not use MPLS GAL. Separate UDP destination port numbers are user-configured for delay and loss measurements from the range specified in

[[RFC8762](#)]. The querier node uses the destination UDP port number following the guidelines specified in [Section 6 in \[RFC6335\]](#). For both Links and end-to-end SR Policies, no PM session for delay or loss measurement is created on the responder node R5 [[RFC6374](#)].

For Performance Measurement, probe query and response messages are sent as following:

- o For Delay Measurement, the probe messages are sent on the congruent path of the data traffic by the querier node, and are used to measure the delay experienced by the actual data traffic flowing on the Links and SR Policies.
- o For Loss Measurement, the probe messages are sent on the congruent path of the data traffic by the querier node, and are used to collect the receive traffic counters for the incoming link or incoming SID where the probe query messages are received at the responder node (incoming link or incoming SID needed since the responder node does not have PM session state present).

The In-Situ Operations, Administration, and Maintenance (IOAM) mechanisms for SR-MPLS defined in [[I-D.gandhi-mpls-ioam-sr](#)] and for SRv6 defined in [[I-D.ali-spring-ioam-srv6](#)] are used to carry PM information such as timestamp in-band as part of the data packets, and are outside the scope of this document.

[3.1.](#) Example Provisioning Model

An example provisioning model described in [[I-D.gandhi-spring-stamp-srpm](#)] is also applicable to the procedures defined in this document.

[4.](#) Probe Query Message

In this document, UDP path is used for Delay and Loss measurements for Links and end-to-end SR Policies for the probe messages defined in [[RFC6374](#)]. The user-configured destination UDP ports (separate UDP ports for different delay and loss message formats) are used for identifying the probe packets.

[4.1.](#) Delay Measurement Probe Query Message

The message content for Delay Measurement for probe query message using UDP header [[RFC0768](#)] is shown in Figure 1. The DM probe query message is sent with user-configured Destination UDP port number for DM. The Destination UDP port can also be used as Source port for two-way delay measurement, since the message has a flag to distinguish between query and response. The DM probe query message

contains the payload format for delay measurement defined in [Section 3.2 of \[RFC6374\]](#).

```

+-----+
| IP Header                                     |
. Source IP Address = Querier IPv4 or IPv6 Address .
. Destination IP Address = Responder IPv4 or IPv6 Address .
. Protocol = UDP .
.
+-----+
| UDP Header                                   |
. Source Port = As chosen by Querier .
. Destination Port = User-configured Port for Delay Measurement.
.
+-----+
| Payload = Message as specified in Section 3.2 of RFC 6374 |
.
+-----+

```

Figure 1: DM Probe Query Message

It is recommended to use the IEEE 1588v2 Precision Time Protocol (PTP) truncated 64-bit timestamp format as a default format as specified in [Appendix A of \[RFC6374\]](#), with hardware support. As an alternative, Network Time Protocol (NTP) timestamp format can also be used [\[RFC6374\]](#).

[4.2.](#) Loss Measurement Probe Query Message

The message content for Loss measurement probe query message using UDP header [\[RFC0768\]](#) is shown in Figure 2. As shown, the LM probe query message is sent with user-configured Destination UDP port number for LM. Separate Destination UDP ports are used for direct-mode and inferred-mode loss measurements. The Destination UDP port can also be used as Source port for two-way loss measurement, since the message has a flag to distinguish between query and response. The LM probe query message contains the payload format for loss measurement defined in [Section 3.1 of \[RFC6374\]](#).


```

+-----+
| IP Header |
. Source IP Address = Querier IPv4 or IPv6 Address .
. Destination IP Address = Responder IPv4 or IPv6 Address .
. Protocol = UDP .
. .
+-----+
| UDP Header |
. Source Port = As chosen by Querier .
. Destination Port = User-configured Port for Loss Measurement .
. .
+-----+
| Payload = Message as specified in Section 3.1 of RFC 6374 |
. .
+-----+

```

Figure 2: LM Probe Query Message

[4.3.](#) Combined Loss/Delay Measurement Probe Query Message

The message content for combined Loss/Delay measurement probe query message using UDP header [[RFC0768](#)] is shown in Figure 3. As shown, the probe query message is sent with user-configured Destination UDP port number for combined LM/DM message format. Separate Destination UDP ports are used for direct-mode and inferred-mode loss measurements. The Destination UDP port can also be used as Source port for two-way loss/delay measurement, since the message has a flag to distinguish between query and response. The probe query message contains the payload format for combined loss/delay measurement defined in [Section 3.3 of \[RFC6374\]](#).

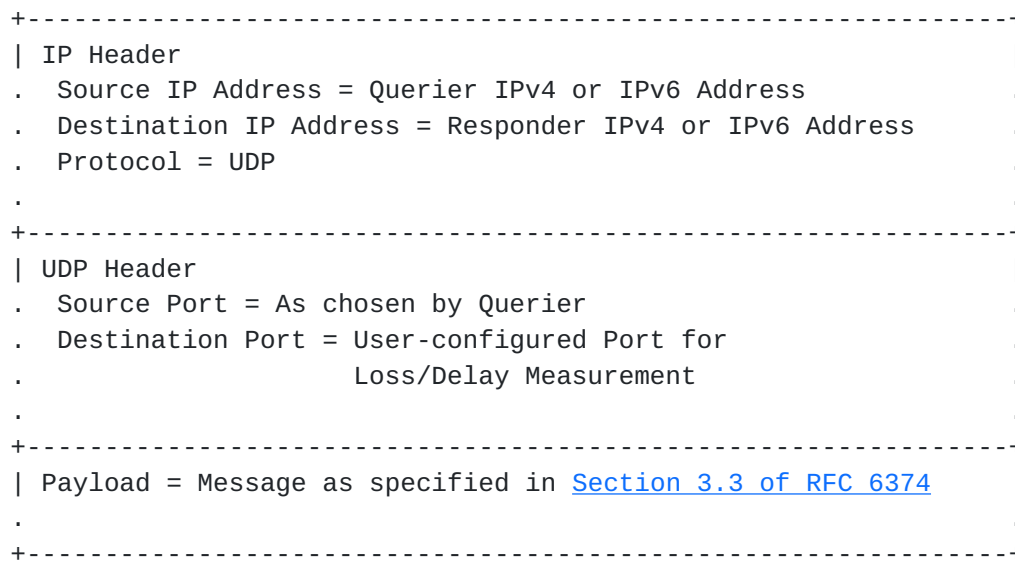


Figure 3: LM/DM Probe Query Message

[4.4.](#) Probe Query Message for Links


The probe query message as defined in Figure 1 is sent on the congruent path of the data traffic for performance Delay measurement. Similarly, the probe query message as defined in Figure 2 is sent on the congruent path of the data traffic for performance Loss measurement.

[4.5.](#) Probe Query Message for SR Policies

The performance delay and loss measurement for segment routing is applicable to both SR-MPLS and SRv6 Policies.

[4.5.1.](#) Probe Query Message for SR-MPLS Policy

The probe query message for end-to-end SR-MPLS Policy performance measurement is sent using its SR-MPLS header containing the MPLS segment list as shown in Figure 4.



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

An SRv6 Policy setup using the SRv6 Segment Routing Header (SRH) and a Segment List is defined in [RFC8754]. For SRv6, network programming is defined in [I-D.ietf-spring-srv6-network-programming]. The probe query messages using UDP header for end-to-end SRv6 Policy performance measurement is sent using its SRv6 Segment Routing Header (SRH) and Segment List as shown in Figure 5.


```

+-----+
| IP Header                                     |
. Source IP Address = Querier IPv6 Address    .
. Destination IP Address = Destination IPv6 Address .
.                                             .
+-----+
| SRH as specified in RFC 8754                 |
. <Segment List>                             .
.                                             .
+-----+
| IP Header                                     |
. Source IP Address = Querier IPv6 Address    .
. Destination IP Address = Responder IPv6 Address .
.                                             .
+-----+
| UDP Header                                   |
. Source Port = As chosen by Querier          .
. Destination Port = User-configured Port     .
.                                             .
+-----+
| Payload = DM Message as specified in Figure 1 |
. Payload = LM Message as specified in Figure 2 .
. Payload = LM/DM Message as specified in Figure 3 .
.                                             .
+-----+

```

Figure 5: Example Probe Query Message for SRv6 Policy

5. Probe Response Message

When the received probe query message does not contain any UDP Return Object (URO) TLV [[RFC7876](#)], the probe response message is sent using the IP/UDP information from the received probe query message. The content of the probe response message is shown in Figure 6.


```

+-----+
| IP Header |
. Source IP Address = Responder IPv4 or IPv6 Address .
. Destination IP Address = Source IP Address from Query .
. Protocol = UDP .
. .
+-----+
| UDP Header |
. Source Port = As chosen by Responder .
. Destination Port = Source Port from Query .
. .
+-----+
| Message as specified in Section 3.2 of RFC 6374 for DM, or |
. Message as specified in Section 3.1 of RFC 6374 for LM, or .
. Message as specified in Section 3.3 of RFC 6374 for LM/DM .
. .
+-----+

```

Figure 6: Probe Response Message

When the received probe query message contains UDP Return Object (URO) TLV [[RFC7876](#)], the probe response message uses the IP/UDP information from the URO in the probe query message. The content of the probe response message is shown in Figure 7.

```

+-----+
| IP Header |
. Source IP Address = Responder IPv4 or IPv6 Address .
. Destination IP Address = URO.Address .
. Protocol = UDP .
. .
+-----+
| UDP Header |
. Source Port = As chosen by Responder .
. Destination Port = URO.UDP-Destination-Port .
. .
+-----+
| Message as specified in Section 3.2 of RFC 6374 for DM, or |
. Message as specified in Section 3.1 of RFC 6374 for LM, or .
. Message as specified in Section 3.3 of RFC 6374 for LM/DM .
. .
+-----+

```

Figure 7: Probe Response Message Using URO from Probe Query

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

[5.1.1.](#) Links and SR Policies

In one-way performance measurement mode, the probe response message as defined in Figure 6 or Figure 7 is sent out-of-band to the querier node, for both Links and SR Policies.

The querier node can receive probe response message back by setting its own IP address as Source Address of the header or by adding URO TLV in the probe query message and setting its own IP address in the IP Address in the URO TLV (Type=131) [[RFC7876](#)]. The "control code" in the probe query message is set to "out-of-band response requested". The "Source Address" TLV (Type 130), and "Return Address" TLV (Type 1), if present in the probe query message, are not used to send probe response message. In this delay measurement mode, as per Reference Topology, timestamps t1 and t2 are collected by the probes to measure one-way delay.

[5.1.2.](#) Probe Response Message to Controller

As shown in the Reference Topology, if the querier node requires the probe response message to be sent to the controller R100, it adds URO TLV in the probe query message and sets the IP address of R100 in the IP Address field and user-configured UDP port for DM and for LM in the UDP-Destination-Port field of the URO TLV (Type=131) [[RFC7876](#)].

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

[5.2.1.](#) Links

In two-way performance measurement mode, when using a bidirectional link, the probe response message as defined in Figure 6 or Figure 7 is sent back on the congruent path of the data traffic to the querier node for Links. In this case, the "control code" in the probe query message is set to "in-band response requested" [[RFC6374](#)]. In this delay measurement mode, as per Reference Topology, timestamps t1, t2, t3 and t4 are collected by the probes to measure two-way delay.

[5.2.2.](#) SR Policies

In two-way performance measurement mode, when using a bidirectional path, the probe response message is sent back on the congruent path of the data traffic to the querier node for end-to-end SR Policies measurements. In this case, the "control code" in the probe query message is set to "in-band response requested" [[RFC6374](#)].

5.2.3. Return Path TLV

For two-way performance measurement, the querier node can request the responder node to send a response message back on a given reverse path (e.g. co-routed path for two-way measurement). Return Path TLV defined in [[I-D.gandhi-mpls-rfc6374-sr](#)] is used to carry reverse SR path information as part of the payload of the probe query message.

Additional Sub-TLVs are defined in this document for the Return Path TLV for the following Types:

- o Type (value TBA1): SRv6 Segment List of the Reverse SR Path
- o Type (value TBA2): SRv6 Binding SID [[I-D.ietf-pce-binding-label-sid](#)] of the Reverse SR Policy

5.2.4. Probe Response Message for SR-MPLS Policy

The message content for sending probe response message on the congruent path of the data traffic for two-way end-to-end SR-MPLS Policy performance measurement is shown in Figure 8. The SR-MPLS label stack in the packet header is built using the Segment List received in the Return Path TLV in the probe query message.

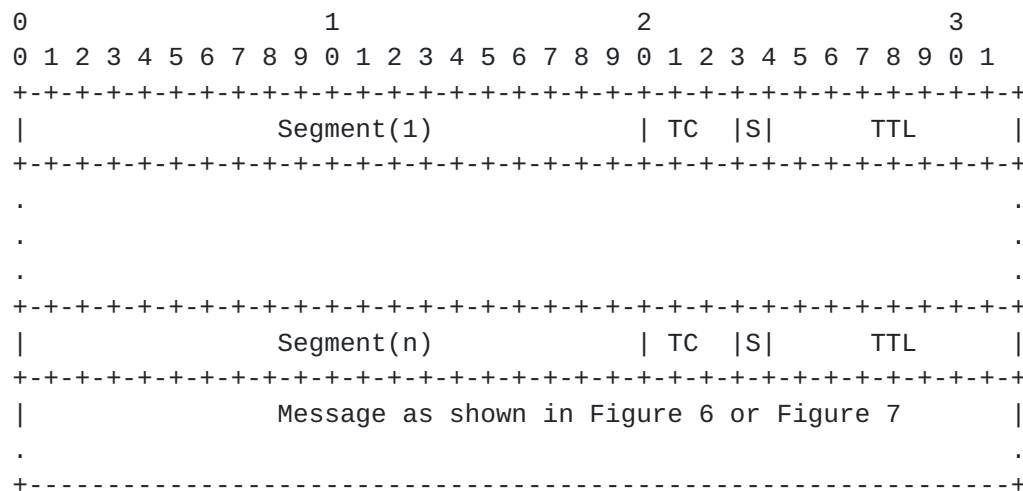


Figure 8: Example Probe Response Message for SR-MPLS Policy

The Path Segment Identifier (PSID)

[[I-D.ietf-spring-mpls-path-segment](#)] of the forward SR-MPLS Policy can be used to find the reverse SR-MPLS Policy to send the probe response message for two-way measurement in the absence of Return Path TLV.

5.2.5. Probe Response Message for SRv6 Policy

The message content for sending probe response message on the congruent path of the data traffic for two-way end-to-end SRv6 Policy performance measurement is shown in Figure 9. For SRv6 Policy using SRH, the SRv6 SID list in the SRH of the probe response message is built using the SRv6 Segment List received in the Return Path TLV in the probe query message.

```

+-----+
| IP Header                                     |
. Source IP Address = Responder IPv6 Address   .
. Destination IP Address = Destination IPv6 Address .
.                                             .
+-----+
| SRH as specified in RFC 8754                 |
. <Segment List>                             .
.                                             .
+-----+
| IP Header                                     |
. Source IP Address = Responder IPv6 Address   .
. Destination IP Address = Querier IPv6 Address .
.                                             .
+-----+
| UDP Header                                   |
. Source Port = As chosen by Responder         .
. Destination Port = Source Port from Query    .
.                                             .
+-----+
| Message as specified in Section 3.2 of RFC 6374 for DM, or |
. Message as specified in Section 3.1 of RFC 6374 for LM, or .
. Message as specified in Section 3.3 of RFC 6374 for LM/DM .
.                                             .
+-----+

```

Figure 9: Example Probe Response Message for SRv6 Policy

5.3. Loopback Measurement Mode

The Loopback measurement mode defined in [Section 2.8 of \[RFC6374\]](#) can be used to measure round-trip delay of a bidirectional SR Path. The IP header of the probe query message contains the destination address equals to the querier node address and the source address equals to the responder address. Optionally, the probe query message can carry the reverse path information (e.g. reverse path label stack for SR-MPLS) as part of the SR header. The responder node does not process the probe messages and generate response messages, and hence Loopback Request object (Type 3) is not required for SR. In this delay

measurement mode, as per Reference Topology, timestamps t1 and t4 are collected by the probes to measure round-trip delay.

6. Performance Measurement for P2MP SR Policies

For P2MP SR Policies [[I-D.voyer-spring-sr-replication-segment](#)], the procedure defined in [[I-D.gandhi-spring-stamp-srpm](#)] is also applicable to the procedures defined in this document (using the [RFC 6374](#) messages in the payload).

7. ECMP Support for SR Policies

For handling ECMP of SR Policies, the procedure defined in [[I-D.gandhi-spring-stamp-srpm](#)] is also applicable to the procedure defined in this document.

8. Additional Probe Message Processing Rules

The additional probe message processing rules defined in [[I-D.gandhi-spring-stamp-srpm](#)] are also applicable to the procedures defined in this document.

9. Sequence Numbers

The message formats for DM and LM [[RFC6374](#)] can carry either timestamp or sequence number but not both. There are case where both timestamp and sequence number are desired for both DM and LM. Sequence numbers can be useful when some probe query messages are lost or they arrive out of order. In addition, the sequence numbers can be useful for detecting denial-of-service (DoS) attacks on UDP ports.

9.1. Sequence Number TLV in Unauthenticated Mode

[RFC6374] defines DM and LM probe query and response messages that can include one or more optional TLVs. New TLV Type (value TBA3) is defined in this document to carry sequence number for probe query and response messages for delay and loss measurement. The format of the Sequence Number TLV in unauthenticated mode is shown in Figure 10.

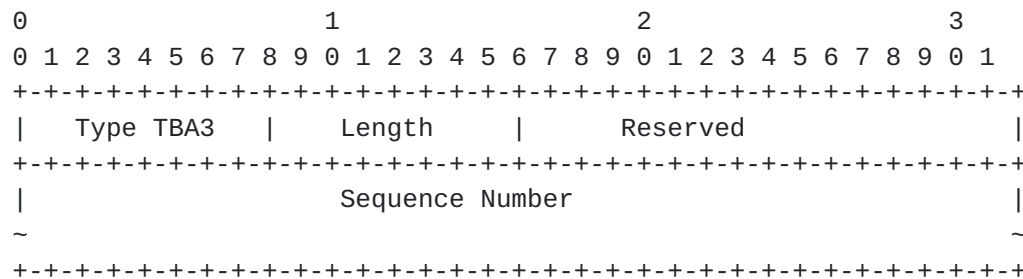


Figure 10: Sequence Number TLV - Unauthenticated Mode

- o The sequence numbers start with 0 and are incremented by one for each subsequent probe query packet.
- o The sequence number are independent for DM and LM messages.
- o The sequence number can be of any length determined by the querier node.
- o The Sequence Number TLV is optional.
- o The querier node SHOULD only insert one Sequence Number TLV in the probe query message and the responder node in the probe response message SHOULD return the first Sequence Number TLV from the probe query message and ignore the other Sequence Number TLVs if present.
- o When Sequence Number TLV is added, the DM and LM messages SHOULD NOT carry sequence number in the timestamp field of the message.

9.2. Sequence Number TLV in Authenticated Mode

The probe query and response packet format in authenticated mode includes a key Hashed Message Authentication Code (HMAC) ([RFC2104](#)) hash. Each probe query and response messages are authenticated by adding Sequence Number with Hashed Message Authentication Code (HMAC) TLV. It can use HMAC-SHA-256 truncated to 128 bits (similarly to the use of it in IPSec defined in [RFC4868](#)); hence the length of the HMAC field is 16 octets.

In authenticated mode, only the sequence number is encrypted, and the other payload fields are sent in clear text. The probe packet MAY include Comp.MBZ (Must Be Zero) variable length field to align the packet on 16 octets boundary.

The computation of HMAC field using HMAC-SHA1 can be used with the procedure defined in this document. HMAC uses own key and the definition of the mechanism to distribute the HMAC key is outside the

scope of this document. Both the authentication type and key can be user-configured on both the querier and responder nodes.

The format of the Sequence Number TLV in authentication mode is shown in Figure 11.

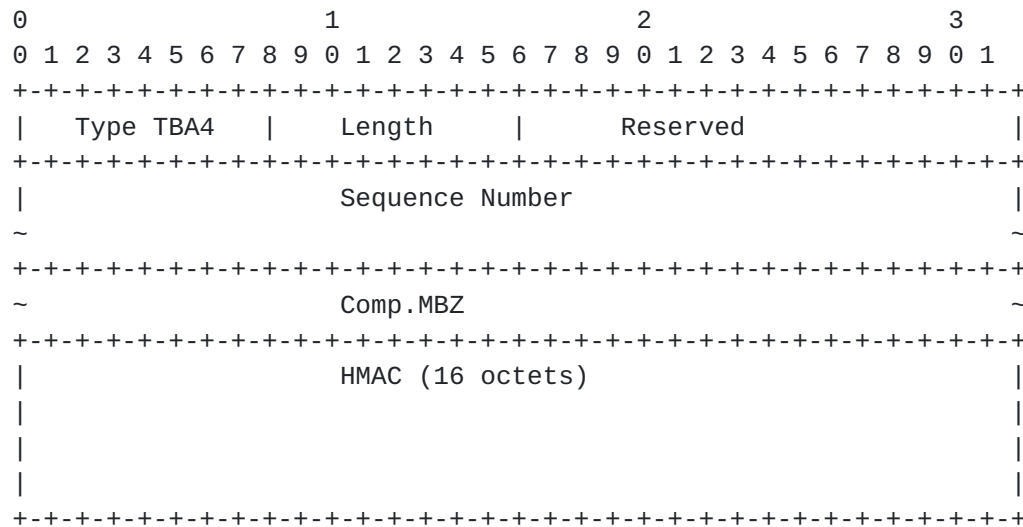


Figure 11: Sequence Number TLV - Authenticated Mode

- o This TLV is mandatory in the authenticated mode.
- o The node MUST discard the probe message if HMAC is invalid.
- o The Sequence Number follows the same processing rule as defined in the unauthenticated mode.

10. Performance Delay and Liveness Monitoring

The procedure defined in this document for delay measurement can also be applied to liveness monitoring of Links and SR Paths. Liveness failure is notified when consecutive N number of probe response messages are not received back at the querier node, where N is locally provisioned value. Note that for one-way and two-way modes, the detection interval and scale for number of sessions need to account for the processing of the probe messages which are punted out of fast path in forwarding (to slow path or control plane), and re-injected on the responder node.

11. 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. The security considerations described in [Section 8 of \[RFC6374\]](#) are applicable to this specification, and particular attention should be paid to the last three paragraphs.

If desired, attacks can be mitigated by performing basic validation and sanity checks, at the querier node, 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 [\[RFC8754\]](#). Hence, probe messages for SRv6 may not need authentication mode. Cryptographic measures may be enhanced by the correct configuration of access-control lists and firewalls.

[12.](#) IANA Considerations

IANA is requested to allocate the values for the following Sub-TLV Types for the Return Path TLV for [RFC 6374](#) from the sub-registry "Return Path Sub-TLV Type" of the "MPLS Loss/Delay Measurement TLV Object" registry contained within the "Generic Associated Channel (G-ACh) Parameters" registry set:

- o Type TBA1: SRv6 Segment List of the Reverse SR Path
- o Type TBA2: SRv6 Binding SID of the Reverse SR Policy

IANA is also requested to allocate the values for the following Sequence Number TLV Types for [RFC 6374](#) to be carried in the probe query and response messages for delay and loss measurement from the "MPLS Loss/Delay Measurement TLV Object" registry contained within the "Generic Associated Channel (G-ACh) Parameters" registry set:

- o Type TBA3: Sequence Number TLV in Unauthenticated Mode
- o Type TBA4: Sequence Number TLV in Authenticated Mode

[13.](#) References

[13.1.](#) 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>>.

- [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>>.
- [RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", [RFC 6374](#), DOI 10.17487/RFC6374, September 2011, <<https://www.rfc-editor.org/info/rfc6374>>.
- [RFC7876] Bryant, S., Sivabalan, S., and S. Soni, "UDP Return Path for Packet Loss and Delay Measurement for MPLS Networks", [RFC 7876](#), DOI 10.17487/RFC7876, July 2016, <<https://www.rfc-editor.org/info/rfc7876>>.
- [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>>.
- [I-D.gandhi-mpls-rfc6374-sr]
Gandhi, R., Filsfils, C., Voyer, D., Salsano, S., and M. Chen, "Performance Measurement Using [RFC 6374](#) for Segment Routing Networks with MPLS Data Plane", [draft-gandhi-mpls-rfc6374-sr-03](#) (work in progress), June 2020.
- [I-D.gandhi-spring-stamp-srpm]
Gandhi, R., Filsfils, C., Voyer, D., Chen, M., and B. Janssens, "Performance Measurement Using STAMP for Segment Routing Networks", [draft-gandhi-spring-stamp-srpm-01](#) (work in progress), June 2020.

13.2. Informative References

- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", [RFC 2104](#), DOI 10.17487/RFC2104, February 1997, <<https://www.rfc-editor.org/info/rfc2104>>.
- [RFC4868] Kelly, S. and S. Frankel, "Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec", [RFC 4868](#), DOI 10.17487/RFC4868, May 2007, <<https://www.rfc-editor.org/info/rfc4868>>.

- [RFC6335] Cotton, M., Eggert, L., Touch, J., Westerlund, M., and S. Cheshire, "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", [BCP 165](#), [RFC 6335](#), DOI 10.17487/RFC6335, August 2011, <<https://www.rfc-editor.org/info/rfc6335>>.
- [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>>.
- [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>>.
- [I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Sivabalan, S., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-07](#) (work in progress), May 2020.
- [I-D.voyer-spring-sr-replication-segment]
Voyer, D., Filsfils, C., Parekh, R., Bidgoli, H., and Z. Zhang, "SR Replication Segment for Multi-point Service Delivery", [draft-voyer-spring-sr-replication-segment-03](#) (work in progress), June 2020.
- [I-D.ietf-pce-binding-label-sid]
Sivabalan, S., Filsfils, C., Tantsura, J., Hardwick, J., Previdi, S., and C. Li, "Carrying Binding Label/Segment-ID in PCE-based Networks.", [draft-ietf-pce-binding-label-sid-02](#) (work in progress), March 2020.
- [I-D.ietf-spring-srv6-network-programming]
Filsfils, C., Camarillo, P., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "SRv6 Network Programming", [draft-ietf-spring-srv6-network-programming-15](#) (work in progress), March 2020.
- [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-02](#) (work in progress), February 2020.

[I-D.gandhi-mpls-ioam-sr]

Gandhi, R., Ali, Z., Filsfils, C., Brockners, F., Wen, B., and V. Kozak, "MPLS Data Plane Encapsulation for In-situ OAM Data", [draft-gandhi-mpls-ioam-sr-02](#) (work in progress), March 2020.

[I-D.ali-spring-ioam-srv6]

Ali, Z., Gandhi, R., Filsfils, C., Brockners, F., Kumar, N., Pignataro, C., Li, C., Chen, M., and G. Dawra, "Segment Routing Header encapsulation for In-situ OAM Data", [draft-ali-spring-ioam-srv6-02](#) (work in progress), November 2019.

Acknowledgments

The authors would like to thank Patrick Khordoc for the discussions on [RFC 6374](#); Nagendra Kumar and Carlos Pignataro for the discussion on SRv6 Performance Measurement. The authors would like to thank Thierry Couture for the discussions on the use-cases for the performance measurement in segment routing networks. The authors would also like to thank Stewart Bryant for the discussion on UDP port allocation for Performance Measurement and Greg Mirsky for providing useful comments and suggestions.

Contributors

Sagar Soni
Cisco Systems, Inc.
Email: sagsoni@cisco.com

Zafar Ali
Cisco Systems, Inc.
Email: zali@cisco.com

Pier Luigi Ventre
CNIT
Italy
Email: pierluigi.ventre@cnit.it

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

Stefano Salsano
Universita di Roma "Tor Vergata"
Italy

Email: stefano.salsano@uniroma2.it

Mach(Guoyi) Chen
Huawei

Email: mach.chen@huawei.com

