

Workgroup: MPLS Working Group
Internet-Draft: draft-ietf-mpls-rfc6374-sr-05
Published: 28 February 2022
Intended Status: Standards Track
Expires: 1 September 2022
Authors: R. Gandhi, Ed. C. Filsfils
 Cisco Systems, Inc. Cisco Systems, Inc.
 D. Voyer S. Salsano
 Bell Canada Universita di Roma "Tor Vergata"
 M. Chen
 Huawei

Performance Measurement Using RFC 6374 for Segment Routing Networks with MPLS Data Plane

Abstract

Segment Routing (SR) leverages the source routing paradigm. RFC 6374 specifies protocol mechanisms to enable the efficient and accurate measurement of packet loss, one-way and two-way delay, as well as related metrics such as delay variation in MPLS networks. This document utilizes these mechanisms for Performance Delay and Loss Measurements in SR networks with MPLS data plane (SR-MPLS), for both SR-MPLS links and end-to-end SR-MPLS paths including Policies. In addition, this document defines Return Path TLV and Block Number TLV extensions for RFC 6374.

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 1 September 2022.

Copyright Notice

Copyright (c) 2022 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Introduction
2.	Conventions Used in This Document
2.1.	Requirements Language
2.2.	Abbreviations
2.3.	Reference Topology
3.	Overview
4.	Query and Response Messages
4.1.	Query Message for SR-MPLS Links and Policies
4.1.1.	Query Message for SR-MPLS Links
4.1.2.	Query Message for SR-MPLS Policies
4.2.	Response Message for SR-MPLS Links and Policies
4.2.1.	One-way Measurement Mode
4.2.2.	Two-way Measurement Mode
4.2.3.	Loopback Measurement Mode
5.	Delay Measurement
5.1.	Delay Measurement Message Format
5.2.	Timestamps
6.	Loss Measurement
6.1.	Loss Measurement Message Format
6.2.	Combined Loss/Delay Measurement Message Format
6.3.	Counters
7.	TLV Extensions
7.1.	Return Path TLV Extension
7.1.1.	Return Path Sub-TLV Extension
7.2.	Block Number TLV Extension
8.	Performance Measurement for P2MP SR-MPLS Policies
9.	ECMP for SR-MPLS Policies
10.	SR-MPLS Link Extended TE Metrics Advertisements
11.	Backwards Compatibility
12.	Security Considerations
13.	IANA Considerations
14.	References
14.1.	Normative References
14.2.	Informative References
	Acknowledgments
	Contributors
	Authors' Addresses

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) 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. A comprehensive SR Performance Measurement (PM) toolset is one of the essential requirements to measure network performance to provide Service Level Agreements (SLAs).

[[RFC6374](#)] specifies protocol mechanisms to enable the efficient and accurate measurement of performance metrics in MPLS networks using query and response messages. [[RFC7876](#)] specifies mechanisms for sending and processing out-of-band responses over an UDP return path when receiving RFC 6374 based query messages. These mechanisms are also well-suited in SR-MPLS networks.

This document utilizes the mechanisms defined in [[RFC6374](#)] for Performance Delay and Loss Measurements in SR-MPLS networks, for both SR-MPLS links and end-to-end SR-MPLS paths including Policies. In addition, this document defines Return Path TLV and Block Number TLV extensions for [[RFC6374](#)].

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

ACH: Associated Channel Header.

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.

SR-MPLS: Segment Routing with MPLS data plane.

TC: Traffic Class.

TE: Traffic Engineering.

URO: UDP Return Object.

2.3. Reference Topology

In the reference topology shown in Figure 1, the querier node Q1 initiates a query message and the responder node R1 transmits a response message for the query message received. The response message may be sent back to the querier node Q1 in-band on the same path (same set of links and nodes) or a different path in the reverse direction.

SR is enabled with MPLS data plane on nodes Q1 and R1. The nodes Q1 and R1 may be directly connected via a link enabled with MPLS (Section 2.9.1 of [[RFC6374](#)]) or a Point-to-Point (P2P) SR-MPLS path [[RFC8402](#)]. The link may be a physical interface, virtual link, or Link Aggregation Group (LAG) [[IEEE802.1AX](#)], or LAG member link. The SR-MPLS path may be an SR-MPLS Policy [[I-D.ietf-spring-segment-routing-policy](#)] on node Q1 (called head-end) with destination to node R1 (called tail-end).




Figure 1: Reference Topology

3. Overview

For delay and loss measurement in SR-MPLS networks, the procedures defined in [RFC6374] are used in this document. Note that the one-way, two-way and round-trip delay measurements are defined in Section 2.4 of [RFC6374] and are further described in this document for SR-MPLS networks. Similarly, the packet loss measurement is defined in Section 2.2 of [RFC6374] and is further described in this document for SR-MPLS networks.

In SR-MPLS networks, the query and response messages defined in [RFC6374] are sent as following:

- *For delay measurement, the query messages MUST be sent in-band (on the same path as data traffic) for SR-MPLS links and end-to-end SR-MPLS paths to collect transmit and receive timestamps.
- *For loss measurement, the query messages MUST be sent in-band (on the same path as data traffic) for SR-MPLS links and end-to-end SR-MPLS paths to collect transmit and receive traffic counters (e.g. for traffic received on the incoming link for the link packet loss and for the incoming Path Segment Identifier (PSID) [I-D.ietf-spring-mpls-path-segment] for the end-to-end SR-MPLS path packet loss).

It may be desired in SR-MPLS networks that the same path (same set of links and nodes) between the querier and responder be used in both directions of the measurement. This is achieved by using the SR-MPLS Return Path TLV extension defined in this document.

The packet loss measurement using Alternate-Marking Method defined in [RFC8321] requires collecting Block Number of the traffic counters. This is achieved by using the Block Number TLV extension defined in this document.

The performance measurement procedure for SR-MPLS links can be used to compute extended Traffic Engineering (TE) metrics for delay and loss as described in this document. The metrics are advertised in the network using the routing protocol extensions defined in [\[RFC7471\]](#), [\[RFC8570\]](#), and [\[RFC8571\]](#).

4. Query and Response Messages

As described in Section 2.9.1 of [\[RFC6374\]](#), the query and response messages flow over an MPLS Generic Associated Channel (G-ACh). These query and response messages contain G-ACh Label (GAL) (value 13, with S=1). The GAL is followed by an Associated Channel Header (ACH), where Channel Type identifies the measurement message type, and the message payload following the ACH as shown in Figure 2.




Figure 2: RFC 6374 Query and Response Message Header

4.1. Query Message for SR-MPLS Links and Policies

4.1.1. Query Message for SR-MPLS Links

A query message as shown in Figure 2 is sent over the SR-MPLS links for both delay and loss measurement using the procedures described in [\[RFC6374\]](#). For SR-MPLS links, the TTL value MUST be set to 255 in the SR-MPLS header. SR-MPLS encapsulation (e.g. using adjacency SID of the link) can be added for transmitting the query messages for SR-MPLS links.

4.1.2. Query Message for SR-MPLS Policies

An SR-MPLS Policy may contain a number of Segment Lists (SLs) (i.e. stack of MPLS labels) [\[I-D.ietf-spring-segment-routing-policy\]](#). The query messages MUST be transmitted for each SL of the SR-MPLS Policy. A query message for an end-to-end SR-MPLS Policy, for both delay and loss measurement, contains an SR-MPLS label stack, with the G-ACh Label (GAL) at the bottom of the stack (with S=1) as shown in Figure 3. For SR-MPLS Policies, the TTL value MUST be set to 255 in the SR-MPLS header.

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Label(1)          | TC  | S |          TTL          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
.
.
.
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Label(n)          | TC  | S |          TTL          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          GAL (value 13)    | TC  | S |          TTL          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| 0 0 0 1 | Version | Reserved |          Channel Type          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 3: Example Query Message Header for an End-to-end SR-MPLS Policy

The SR-MPLS label stack can be empty (format as shown in Figure 2) to indicate the Implicit NULL label case.

For an SR-MPLS Policy, in order to ensure that the query message is processed by the intended responder, Destination Address TLV (Type 129) [[RFC6374](#)] containing the address of the responder can be sent in the query message. The responder that supports this TLV MUST return Success in "Control Code" [[RFC6374](#)] if it is the intended destination for the query. Otherwise, it MUST return 0x15: Error - Invalid Destination Node Identifier [[RFC6374](#)]. The Destination Address TLV is applicable to P2P SR-MPLS Policy only.

4.2. Response Message for SR-MPLS Links and Policies

4.2.1. One-way Measurement Mode

In one-way measurement mode defined in Section 2.4 of [[RFC6374](#)], the querier can receive "out-of-band" response messages with IP/UDP header by properly setting the UDP Return Object (URO) TLV in the query message. The URO TLV (Type=131) is defined in [[RFC7876](#)] and includes the UDP-Destination-Port and IP Address. When the querier sets an IP address and an UDP port in the URO TLV, the response message MUST be sent to that IP address as the destination address and UDP port as the destination port. In addition, the "Control Code" in the query message MUST be set to "out-of-band response requested" [[RFC6374](#)].

In one-way delay measurement mode, as per Reference Topology in Figure 1, the timestamps T1 and T2 are collected by the query and response messages. Both these timestamps are used to measure one-way delay as (T2 - T1).

4.2.2. Two-way Measurement Mode

In two-way measurement mode defined in Section 2.4 of [RFC6374], the response messages are preferred to be sent back in-band on the same link or the same end-to-end SR-MPLS path (same set of links and nodes) in the reverse direction to the querier.

For SR-MPLS links, the response message (format as shown in Figure 2) is sent back on the same incoming link where the query message is received. In this case, the "Control Code" in the query message MUST be set to "in-band response requested" [RFC6374].

For end-to-end SR-MPLS paths, the responder transmits the response message (example as shown in Figure 3) on a specific return SR-MPLS path [I-D.ietf-pce-sr-bidir-path]. The querier can request in the query message to the responder to send the response message back on a given return path using the SR-MPLS Segment List sub-TLV in the Return Path TLV defined in this document.

In two-way delay measurement mode, as per Reference Topology in Figure 1, all four timestamps T1, T2, T3, and T4 are collected by the query and response messages. All four timestamps are used to measure two-way delay as $((T4 - T1) - (T3 - T2))$.

4.2.3. Loopback Measurement Mode

The Loopback measurement mode defined in Section 2.8 of [RFC6374] is used to measure round-trip delay for a bidirectional circular SR-MPLS path [I-D.ietf-pce-sr-bidir-path]. In this mode, the received query messages MUST NOT be punted out of the fast path in forwarding (i.e. to slow path or control-plane) at the responder. In other words, the responder does not process the payload and generate response messages. The loopback function simply returns the received query message to the querier without responder modifications [RFC6374].

The loopback mode is done by generating "queries" with the Response flag set to 1 and adding the Loopback Request object (Type 3) [RFC6374]. The label stack in query messages in this case carry both the forward and reverse paths in the MPLS header. The GAL is still carried at the bottom of the label stack (with S=1) (example as shown in Figure 3).

In loopback delay measurement mode, as per Reference Topology in Figure 1, the timestamps T1 and T4 are collected by the query messages. Both these timestamps are used to measure round-trip delay as $(T4 - T1)$. In this mode, the round-trip delay includes the processing delay on the responder. The responder processing delay component includes only the time required to loop the test packet

from the incoming interface to the outgoing interface in forwarding plane.

5. Delay Measurement

5.1. Delay Measurement Message Format

As defined in [\[RFC6374\]](#), MPLS DM query and response messages use Associated Channel Header (ACH) (value 0x000C for delay measurement) [\[RFC6374\]](#), which identifies the message type, and the message payload following the ACH. For both SR-MPLS links and end-to-end SR-MPLS Policies, the same MPLS DM ACH value MUST be used.

The DM message payload as defined in Section 3.2 of [\[RFC6374\]](#) is used for delay measurement, for both SR-MPLS links and end-to-end SR-MPLS Policies.

5.2. Timestamps

The Section 3.4 of [\[RFC6374\]](#) defines timestamp format that can be used for delay measurement. The IEEE 1588 Precision Time Protocol (PTP) timestamp format [\[IEEE1588\]](#) is used by default as described in Appendix A of [\[RFC6374\]](#).

6. Loss Measurement

The LM protocol can perform two distinct kinds of loss measurement as described in Section 2.9.8 of [\[RFC6374\]](#).

*In inferred mode, LM will measure the loss of specially generated test messages in order to infer the approximate data plane loss level. Inferred mode LM provides only approximate loss accounting.

*In direct mode, LM will directly measure data plane packet loss. Direct mode LM provides perfect loss accounting, but may require hardware support.

6.1. Loss Measurement Message Format

As defined in [\[RFC6374\]](#), MPLS LM query and response messages use Associated Channel Header (ACH) (value 0x000A for direct loss measurement or value 0x000B for inferred loss measurement), which identifies the message type, and the message payload following the ACH. For both SR-MPLS links and end-to-end SR-MPLS Policies, the same MPLS LM ACH value MUST be used.

The LM message payload as defined in Section 3.1 of [\[RFC6374\]](#) is used for loss measurement, for both SR-MPLS links and end-to-end SR-MPLS Policies.

6.2. Combined Loss/Delay Measurement Message Format

As defined in [RFC6374], Combined DM+LM query and response messages use Associated Channel Header (ACH) (value 0x000D for direct loss and delay measurement or value 0x000E for inferred loss and delay measurement), which identifies the message type, and the message payload following the ACH. For both SR-MPLS links and end-to-end SR-MPLS Policies, the same MPLS DM+LM ACH value MUST be used.

The message payload as defined in Section 3.3 of [RFC6374] is used for combined delay and loss measurement, for both SR-MPLS links and end-to-end SR-MPLS Policies.

6.3. Counters

The Path Segment Identifier (PSID) [I-D.ietf-spring-mpls-path-segment] carried in the received data packet for the traffic flow under measurement can be used for accounting received traffic on the egress node of the SR-MPLS Policy. In direct mode, the PSID in the received query message as shown in Figure 4 can be used to associate the receive traffic counter on the responder to detect the transmit packet loss for the end-to-end SR-MPLS Policy.

In inferred mode, the PSID in the received query messages as shown in Figure 4 can be used to count the received query messages on the responder to detect the transmit packet loss for an end-to-end SR-MPLS Policy.

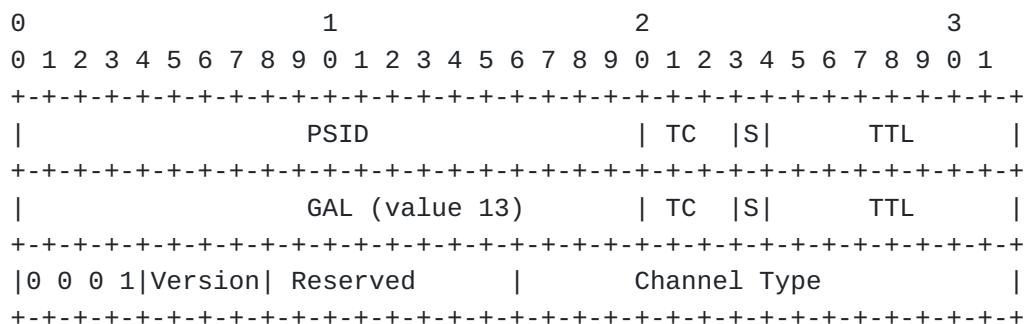


Figure 4: Example With Path Segment Identifier for SR-MPLS Policy

Different values of PSID can be used to measure packet loss per SR-MPLS Policy, per Candidate Path or per Segment List of the SR-MPLS Policy [I-D.ietf-spring-segment-routing-policy].

7. TLV Extensions

7.1. Return Path TLV Extension

In two-way measurement mode, the responder transmits the response message on a specific return path. The querier can request in the query message to the responder to send a response message back on a given return path (e.g. co-routed SR-MPLS path for two-way measurement). This way the responder avoids creating and maintaining extra dynamic SR states for the return paths for two-way measurement.

The querier may not be reachable from the responder. The querier in this case **MUST** send its reachability path information to the responder using the Return Path TLV.

[[RFC6374](#)] defines query and response messages those can include or more optional TLVs. New TLV Type (TBA1) is defined in this document for Return Path TLV to carry return path information in query messages. The format of the Return Path TLV is shown in Figure 5:

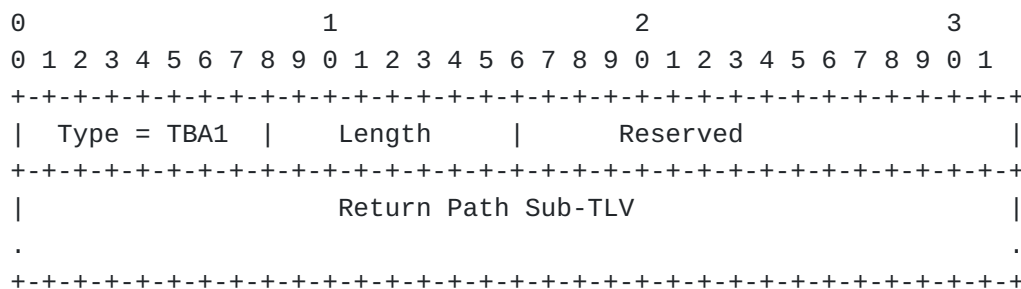


Figure 5: Return Path TLV

The Return Path TLV is a Mandatory TLV Type. The querier **MUST** only insert one Return Path TLV in the query message. The responder that supports this TLV, **MUST** only process the first Return Path TLV and ignore the other Return Path TLVs if present. The responder that supports this TLV, also **MUST** send response message back on the return path specified in the Return Path TLV. The responder also **MUST NOT** add Return Path TLV in the response message. The Reserved field **MUST** be set to 0 and **MUST** be ignored on the receive side.

7.1.1. Return Path Sub-TLV Extension

The Return Path TLV contains a Sub-TLV to carry the return path. The format of the SR-MPLS Segment List Sub-TLV is shown in Figure 6. The SR-MPLS Segment List Sub-TLV contains SR-MPLS Label Stack. The Label entries in the Segment List **MUST** be in network order. The SR-MPLS

Segment List Sub-TLV in the Return Path TLV is of the following Type:

*Type (value 1): SR-MPLS Segment List of the Return Path



Figure 6: SR-MPLS Segment List Sub-TLV in Return Path TLV

An SR-MPLS Segment List Sub-TLV may carry only Binding SID [[I-D.ietf-pce-binding-label-sid](#)] of the Return SR-MPLS Policy.

The Return Path TLV MUST carry only one Return Path Sub-TLV. The responder that supports this Sub-TLV, MUST only process the first Return Path Sub-TLV and ignore the other Return Path Sub-TLVs if present. The responder that supports this Sub-TLV, also MUST send response message back on the return path specified in the Return Path Sub-TLV. The Reserved field MUST be set to 0 and MUST be ignored on the receive side.

Note that in addition to the P2P SR-MPLS paths, the SR-MPLS Segment List Sub-TLV is also applicable to the P2MP SR-MPLS paths. For example, for P2MP SR-MPLS paths, it may only carry the Node Segment Identifier of the querier in order for the response to follow an SR-MPLS path back to the querier.

7.2. Block Number TLV Extension

The direct mode loss measurement using Alternate-Marking Method defined in [[RFC8321](#)] requires collecting Block Number of the counters for the data traffic flow under measurement. To be able to correlate the transmit and receive traffic counters of the matching Block Number, the Block Number of the traffic counters is carried in the LM query and response messages.

[RFC6374] defines query and response messages those can include one or more optional TLVs. New TLV Type (value TBA2) is defined in this document to carry the Block Number (8-bit) of the traffic counters in the LM query and response messages. The format of the Block Number TLV is shown in Figure 7:




Figure 7: Block Number TLV

The Block Number TLV is a Mandatory TLV Type. The querier MUST only insert one Block Number TLV in the query message to identify the Block Number for the traffic counters in the forward direction. The responder that supports this TLV, MUST only insert one Block Number TLV in the response message to identify the Block Number for the traffic counters in the reverse direction. The responder also MUST return the first Block Number TLV from the query message and ignore the other Block Number TLVs if present. The R Flag MUST be clear in the query message and set in the response message. The Reserved field MUST be set to 0 and MUST be ignored on the receive side.

8. Performance Measurement for P2MP SR-MPLS Policies

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

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

*The querier root node transmits query messages using the Tree-SID defined in [\[I-D.ietf-pim-sr-p2mp-policy\]](#) for the P2MP SR-MPLS Policy as shown in Figure 8. The query messages may contain the replication SID as defined in [\[I-D.ietf-spring-sr-replication-segment\]](#).

*Each responder leaf node MUST send its node address in the "Source Address" TLV (Type 130) [[RFC6374](#)] in the response messages. This TLV allows the querier root node to identify the responder leaf nodes of the P2MP SR-MPLS Policy.

*The P2MP root node measures the delay and loss performance for each P2MP leaf node individually.

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     | TC  |S|          TTL          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
.
.
.
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          GAL (value 13)          | TC  |S|          TTL          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0 0 0 1|Version| Reserved          |          Channel Type          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 8: Example Query with Tree-SID for SR-MPLS Policy

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

9. ECMP for SR-MPLS Policies

An SR-MPLS 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-MPLS Policy can result in ECMP paths via transit nodes part of that Anycast group. The query and response messages SHOULD be sent to traverse different ECMP paths to measure delay of each of the ECMP path of an SR-MPLS Policy.

Forwarding plane has various hashing functions available to forward packets on specific ECMP paths. For SR-MPLS Policy, sweeping of entropy label [[RFC6790](#)] values can be used in query and response messages to take advantage of the hashing function in forwarding plane to influence the ECMP path taken by them.

The considerations for loss measurement for different ECMP paths of an SR-MPLS Policy are outside the scope of this document.

10. SR-MPLS Link Extended TE Metrics Advertisements

The extended TE metrics for SR-MPLS link delay and loss can be computed using the performance measurement procedures described in this document to advertise in the routing domain as follows:

- *For OSPF, ISIS, and BGP-LS, protocol extensions defined in [\[RFC7471\]](#), [\[RFC8570\]](#), and [\[RFC8571\]](#), respectively, are used for advertising the extended TE link delay and loss metrics in the network.
- *The extended TE link delay and loss metrics are advertised for Layer-2 LAG bundle member links in OSPF [[I-D.ietf-lsr-ospf-12bundles](#)] and ISIS [[RFC8668](#)] using the same protocol extensions defined in [\[RFC7471\]](#) and [\[RFC8570\]](#), respectively.
- *The advertised delay-variance metric, Packet Delay Variation (PDV) is computed as described in Section 4.2 of [\[RFC5481\]](#).
- *In the absence of one-way delay measurement, the extended TE link one-way delay metrics can be computed using the two-way and round-trip delay values by dividing the values by 2.

11. Backwards Compatibility

The procedures defined in this document are backwards compatible with the procedures defined in [\[RFC6374\]](#) at both querier and responder. If the responder does not support the new Mandatory TLV Types defined in this document, it MUST return Error 0x17: Unsupported Mandatory TLV Object as per [\[RFC6374\]](#).

12. Security Considerations

This document describes the procedures for performance delay and loss measurement for SR-MPLS networks, for both SR-MPLS links and end-to-end SR-MPLS Policies using the mechanisms defined in [\[RFC6374\]](#) and [\[RFC7876\]](#). The security considerations covered in [\[RFC6374\]](#), [\[RFC7471\]](#), [\[RFC8570\]](#), [\[RFC8571\]](#), and [\[RFC7876\]](#) also apply to the procedures in this document.

The procedure defined in this document is intended for deployment in limited domains [\[RFC8799\]](#). As such, it assumes that a querier node involved in the measurement operation has previously verified the integrity of the path and the identity of the far-end responder.

If desired, attacks can be mitigated by performing basic validation and sanity checks, at the querier, of the timestamp and counter fields in received 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 test cycle.

The extensions defined in this document may be used for potential "proxying" attacks. For example, a querier may specify a return path that has a destination different from that of the responder. But normally, such attacks will not happen in an SR-MPLS domain where the queriers and responders belong to the same domain [[RFC8799](#)]. In order to prevent using the extension defined in this document for proxying any possible attacks, the return path has destination to the same node where the forward path is from. The responder may drop the query messages when it cannot determine whether the Return Path has the destination to the querier. That means, the querier may send a proper source address in the "Source Address" TLV according to the specified Return Path to help the responder to make that decision.

13. IANA Considerations

IANA is requested to allocate values for the following Mandatory TLV Types for [[RFC6374](#)] from the "MPLS Loss/Delay Measurement TLV Object" registry contained within the "Generic Associated Channel (G-ACh) Parameters" registry set:

Value	Description	Reference
TBA1	Return Path TLV	This document
TBA2	Block Number TLV	This document

Table 1: MPLS Loss/Delay Measurement TLV Types

The Block Number TLV is carried in the query and response messages and Return Path TLV is carried in the query messages.

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
1 - 175	IETF Review	This document
176 - 239	First Come First Served	This document
240 - 251	Experimental Use	This document
252 - 254	Private Use	This document

Table 2: Return Path Sub-TLV Type Registry

This document defines the following values in the Return Path Sub-TLV Type sub-registry:

Value	Description	Reference
0	Reserved	This document
1	SR-MPLS Segment List of the Return Path	This document
255	Reserved	This document

Table 3: Return Path Sub-TLV Types

14. References

14.1. 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>>.
- [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>>.

14.2. Informative References

- [IEEE1588] IEEE, "1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", March 2008.
- [RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", RFC 5481, DOI 10.17487/RFC5481, March 2009, <<https://www.rfc-editor.org/info/rfc5481>>.
- [RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", RFC 6790, DOI 10.17487/RFC6790, November 2012, <<https://www.rfc-editor.org/info/rfc6790>>.
- [RFC7471] Giacalone, S., Ward, D., Drake, J., Atlas, A., and S. Previdi, "OSPF Traffic Engineering (TE) Metric

Extensions", RFC 7471, DOI 10.17487/RFC7471, March 2015, <<https://www.rfc-editor.org/info/rfc7471>>.

[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>>.

[RFC8321] Fioccola, G., Ed., Capello, A., Cociglio, M., Castaldelli, L., Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi, "Alternate-Marking Method for Passive and Hybrid Performance Monitoring", RFC 8321, DOI 10.17487/RFC8321, January 2018, <<https://www.rfc-editor.org/info/rfc8321>>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shaker, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

[RFC8570] Ginsberg, L., Ed., Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", RFC 8570, DOI 10.17487/RFC8570, March 2019, <<https://www.rfc-editor.org/info/rfc8570>>.

[RFC8571] Ginsberg, L., Ed., Previdi, S., Wu, Q., Tantsura, J., and C. Filsfils, "BGP - Link State (BGP-LS) Advertisement of IGP Traffic Engineering Performance Metric Extensions", RFC 8571, DOI 10.17487/RFC8571, March 2019, <<https://www.rfc-editor.org/info/rfc8571>>.

[RFC8668] Ginsberg, L., Ed., Bashandy, A., Filsfils, C., Nanduri, M., and E. Aries, "Advertising Layer 2 Bundle Member Link Attributes in IS-IS", RFC 8668, DOI 10.17487/RFC8668, December 2019, <<https://www.rfc-editor.org/info/rfc8668>>.

[RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <<https://www.rfc-editor.org/info/rfc8799>>.

[I-D.ietf-spring-segment-routing-policy]

Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", Work in Progress, Internet-Draft, draft-ietf-spring-segment-routing-policy-18, 17 February 2022, <<https://www.ietf.org/archive/id/draft-ietf-spring-segment-routing-policy-18.txt>>.

[I-D.ietf-pim-sr-p2mp-policy] (editor), D. V., Filsfils, C., Parekh, R., Bidgoli, H., and Z. Zhang, "Segment Routing Point-to-

Multipoint Policy", Work in Progress, Internet-Draft, draft-ietf-pim-sr-p2mp-policy-03, 23 August 2021, <<https://www.ietf.org/archive/id/draft-ietf-pim-sr-p2mp-policy-03.txt>>.

[I-D.ietf-spring-sr-replication-segment]

(editor), D. V., Filsfils, C., Parekh, R., Bidgoli, H., and Z. Zhang, "SR Replication Segment for Multi-point Service Delivery", Work in Progress, Internet-Draft, draft-ietf-spring-sr-replication-segment-06, 25 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-spring-sr-replication-segment-06.txt>>.

[I-D.ietf-pce-binding-label-sid]

Sivabalan, S., Filsfils, C., Tantsura, J., Previdi, S., and C. L. (editor), "Carrying Binding Label/Segment Identifier (SID) in PCE-based Networks.", Work in Progress, Internet-Draft, draft-ietf-pce-binding-label-sid-13, 10 February 2022, <<https://www.ietf.org/archive/id/draft-ietf-pce-binding-label-sid-13.txt>>.

[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", Work in Progress, Internet-Draft, draft-ietf-spring-mpls-path-segment-07, 20 December 2021, <<https://www.ietf.org/archive/id/draft-ietf-spring-mpls-path-segment-07.txt>>.

[I-D.ietf-lsr-ospf-l2bundles] Talaulikar, K. and P. Psenak, "Advertising L2 Bundle Member Link Attributes in OSPF", Work in Progress, Internet-Draft, draft-ietf-lsr-ospf-l2bundles-02, 22 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-lsr-ospf-l2bundles-02.txt>>.

[I-D.ietf-pce-sr-bidir-path] Li, C., Chen, M., Cheng, W., Gandhi, R., and Q. Xiong, "Path Computation Element Communication Protocol (PCEP) Extensions for Associated Bidirectional Segment Routing (SR) Paths", Work in Progress, Internet-Draft, draft-ietf-pce-sr-bidir-path-08, 9 September 2021, <<https://www.ietf.org/archive/id/draft-ietf-pce-sr-bidir-path-08.txt>>.

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

Acknowledgments

The authors would like to thank Thierry Couture for the discussions on the use-cases for the performance measurement in segment routing

networks. Authors would like to thank Patrick Khordoc for implementing the mechanisms defined in this document. The authors would like to thank Greg Mirsky for providing many useful comments and suggestions. The authors would also like to thank Stewart Bryant, Sam Aldrin, Tarek Saad, and Rajiv Asati for their review comments. Thanks to Huaimo Chen, Yimin Shen, and Xufeng Liu for MPLS-RT expert review.

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