
Stream: Internet Engineering Task Force (IETF)

RFC: [9779](#)

Category: Standards Track

Published: April 2025

ISSN: 2070-1721

Authors:

R. Gandhi, Ed.

C. Filsfils

D. Voyer

S. Salsano

Cisco Systems, Inc.

Cisco Systems, Inc.

Bell Canada

Universita di Roma "Tor Vergata"

M. Chen

Huawei

RFC 9779

Performance Measurement for Segment Routing Networks with the MPLS Data Plane

Abstract

This document specifies the application of the MPLS loss and delay measurement techniques (originally defined in RFCs 6374, 7876, and 9341) within Segment Routing (SR) networks that utilize the MPLS data plane, also referred to as Segment Routing over MPLS (SR-MPLS). SR enables the forwarding of packets through an ordered list of instructions, known as segments, which are imposed at the ingress node. This document defines the procedures and extensions necessary to perform accurate measurement of packet loss and delay in SR-MPLS environments, ensuring that network operators can effectively measure and maintain the quality of service across their SR-based MPLS networks. This includes coverage of links and end-to-end SR-MPLS paths, as well as SR Policies.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9779>.

Copyright Notice

Copyright (c) 2025 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	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
4. Query and Response Messages	6
4.1. Query Message for Links and SR-MPLS Policies	6
4.1.1. Query Message for Links	6
4.1.2. Query Message for SR-MPLS Policies	6
4.2. Response Message for Links and SR-MPLS Policies	7
4.2.1. One-Way Measurement Mode	7
4.2.2. Two-Way Measurement Mode	7
4.2.3. Loopback Measurement Mode	8
5. Delay and Loss Measurement	9
5.1. Delay Measurement Message	9
5.2. Loss Measurement Message	9
5.3. Combined Loss/Delay Measurement Message	9
5.4. Counters	9
5.5. Block Number for Counters	10
6. TLV Extensions	11
6.1. Return Path TLV Extension	11
6.1.1. Return Path Sub-TLV Extension	12
6.2. Block Number TLV Extension	13

7. ECMP for SR-MPLS Policies	13
8. Extended TE Link Metrics Advertisement	14
9. Backwards Compatibility	14
10. Manageability Considerations	14
11. Security Considerations	14
12. IANA Considerations	15
13. References	16
13.1. Normative References	16
13.2. Informative References	16
Acknowledgments	18
Contributors	18
Authors' Addresses	18

1. Introduction

Segment Routing (SR), as specified in [RFC8402], leverages the source routing paradigm and applies to both the Multiprotocol Label Switching (MPLS) and Internet Protocol version 6 (IPv6) data planes. These are referred to as Segment Routing over MPLS (SR-MPLS) and Segment Routing over IPv6 (SRv6), respectively. SR takes advantage of Equal-Cost Multipaths (ECMPs) between source and transit nodes, between transit nodes, and between transit and destination nodes. SR Policies, defined in [RFC9256], are used to steer traffic through specific, user-defined paths using a list of segments.

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

[RFC6374] specifies protocol mechanisms to enable 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.

[RFC7876] specifies mechanisms for sending and processing out-of-band responses over a UDP return path when receiving query messages defined in [RFC6374]. These mechanisms can be applied to SR-MPLS networks.

[RFC9341] defines the Alternate-Marking Method using Block Numbers as a data correlation mechanism for packet loss measurement.

This document utilizes the mechanisms from [RFC6374], [RFC7876], and [RFC9341] for delay and loss measurements in SR-MPLS networks. This includes coverage of links and end-to-end SR-MPLS paths, as well as SR Policies.

This document extends [RFC6374] by defining Return Path and Block Number TLVs (see [Section 6](#)) for delay and loss measurement in SR-MPLS networks. These TLVs can also be used in MPLS Label Switched Paths (LSPs) [RFC3031]. However, the procedure for delay and loss measurement of MPLS LSPs is outside the scope of this document.

2. Conventions Used in This Document

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2.2. Abbreviations

ACH: Associated Channel Header
DM: Delay Measurement
ECMP: Equal-Cost Multipath
G-ACh: Generic Associated Channel
GAL: Generic Associated Channel Label
LM: Loss Measurement
LSE: Label Stack Entry
MPLS: Multiprotocol Label Switching
PSID: Path Segment Identifier
SID: Segment Identifier
SL: Segment List
SR: Segment Routing
SR-MPLS: Segment Routing over MPLS
TC: Traffic Class
TE: Traffic Engineering
TTL: Time to Live

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 on the same path (the same set of links and nodes) or on a different path in the reverse direction from the path taken towards the responder R1.

T1 is a transmit timestamp, and T4 is a receive timestamp; both are added by node Q1. T2 is a receive timestamp, and T3 is a transmit timestamp; both are added by node R1.

SR is enabled with the MPLS data plane on nodes Q1 and R1. The nodes Q1 and R1 are connected via a channel ([Section 2.9.1](#) of [[RFC6374](#)]). The channel may be a directly connected link enabled with MPLS ([Section 2.9.1](#) of [[RFC6374](#)]) or an SR-MPLS path [[RFC8402](#)]. The link may be a physical interface, a virtual link, a Link Aggregation Group (LAG) [[IEEE802.1AX](#)], or a LAG member link. The SR-MPLS path may be an SR-MPLS Policy [[RFC9256](#)] on node Q1 (called the "head-end") with the destination to node R1 (called the "tail-end").

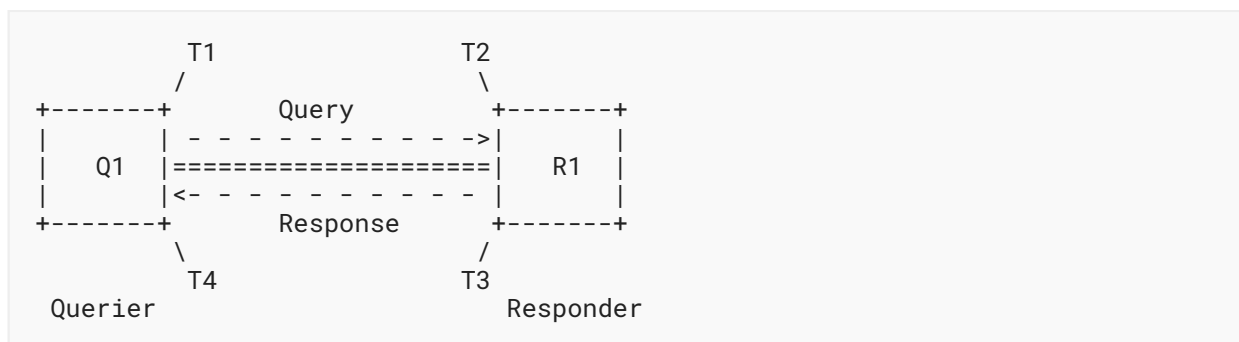


Figure 1: Reference Topology

3. Overview

In this document, the procedures defined in [[RFC6374](#)], [[RFC7876](#)], and [[RFC9341](#)] are utilized for delay and loss measurement in SR-MPLS networks. Specifically, the one-way, two-way, and round-trip delay measurements described in [Section 2.4](#) of [[RFC6374](#)] are further elaborated for application within SR-MPLS networks. Similarly, the packet loss measurement procedures outlined in [Section 2.2](#) of [[RFC6374](#)] are extended for use in SR-MPLS networks.

Packet loss measurement using the Alternate-Marking Method defined in [[RFC9341](#)] may employ the Block Number for data correlation. This is achieved by utilizing the Block Number TLV extension defined in this document.

In SR-MPLS networks, the query messages defined in [RFC6374] **MUST** be transmitted along the same path as the data traffic for links and end-to-end SR-MPLS paths. This is to collect both transmit and receive timestamps for delay measurement and to collect both transmit and receive traffic counters for loss measurement.

If it is desired in SR-MPLS networks that the same path (i.e., the same set of links and nodes) between the querier and responder be used in both directions of the measurement, then this can be achieved by using the Return Path TLV extension defined in this document.

The performance measurement procedures for links can be used to compute extended Traffic Engineering (TE) metrics for delay and loss, as described herein. These metrics are advertised in the network using the routing protocol extensions defined in [RFC7471], [RFC8570], and [RFC8571].

4. Query and Response Messages

4.1. Query Message for Links and SR-MPLS Policies

4.1.1. Query Message for Links

The query message, as defined in [RFC6374], is sent over the links for both delay and loss measurement. In each Label Stack Entry (LSE) [RFC3032] in the MPLS label stack, the TTL value **MUST** be set to 255 [RFC5082].

4.1.2. Query Message for SR-MPLS Policies

An SR-MPLS Policy Candidate-Path may contain a number of Segment Lists (SLs) (i.e., a stack of MPLS labels) [RFC9256]. For delay and/or loss measurement for an end-to-end SR-MPLS Policy, the query messages **MUST** be transmitted for every SL of the SR-MPLS Policy Candidate-Path. This is done by creating a separate session for each SL. Each query message of a session contains an SR-MPLS label stack of the Candidate-Path, with the G-ACh Label (GAL) at the bottom of the stack (with S=1) as shown in Figure 2. In each LSE in the MPLS label stack, the TTL value **MUST** be set to 255 [RFC5082].

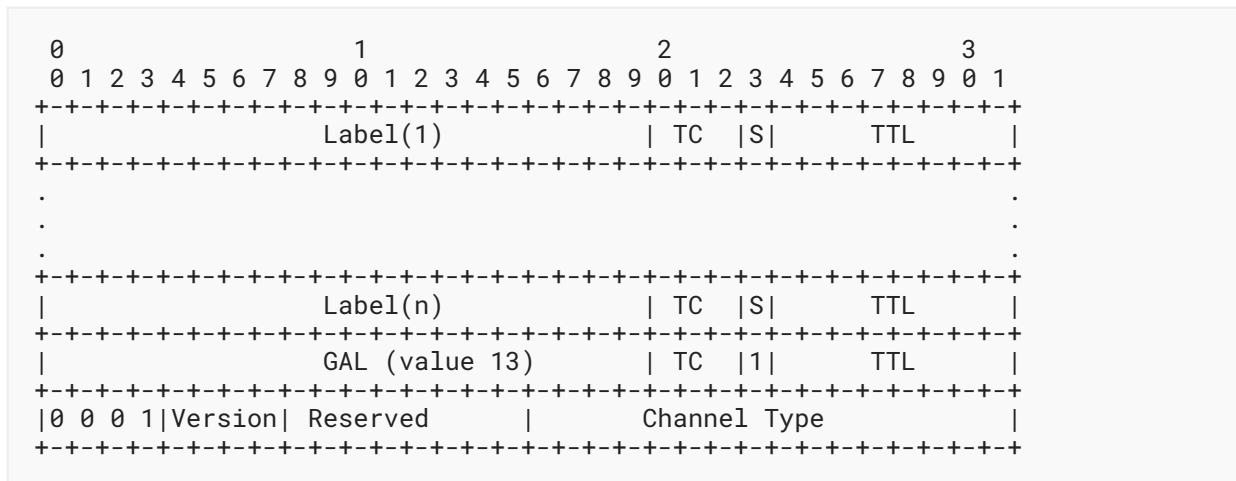


Figure 2: Example Query Message Header for an End-to-End SR-MPLS Policy

The fields 0001, Version, Reserved, and Channel Type shown in Figure 2 are specified in [RFC5586].

The SR-MPLS label stack can be empty in the case of a one-hop SR-MPLS Policy with an Implicit NULL label.

For an SR-MPLS Policy, to ensure that the query message is processed by the intended responder, the Destination Address TLV (Type 129) [RFC6374] containing an address of the responder can be sent in the query messages. The responder that supports this TLV **MUST** return Control Code 0x1 (Success) [RFC6374] if it is the intended destination for the query. Otherwise, it **MUST** return Error 0x15: Invalid Destination Node Identifier [RFC6374].

4.2. Response Message for Links and SR-MPLS Policies

4.2.1. One-Way Measurement Mode

In the one-way measurement mode defined in Section 2.4 of [RFC6374], the querier can receive response messages with an IP/UDP header "out-of-band" 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 a UDP port in the URO TLV, the response message **MUST** be sent to that IP address, with that IP address as the destination address and the 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].

4.2.2. Two-Way Measurement Mode

In the two-way measurement mode defined in Section 2.4 of [RFC6374], the response messages **SHOULD** be sent back one of two ways: either they are sent back in-band on the same link, or they are sent back on the same end-to-end SR-MPLS path (i.e., the same set of links and nodes) in the reverse direction to the querier. This is done in order to perform accurate two-way delay measurement.

For links, the response message as defined in [RFC6374] 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 (see the example as shown in Figure 2) on a specific return SR-MPLS path. In the query message, the querier can request that the responder send the response message back on a given return path using the MPLS Label Stack Sub-TLV in the Return Path TLV defined in this document.

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 path [RFC6374] in SR-MPLS networks. In this mode for SR-MPLS, the received query messages are not punted out of the fast path in forwarding (i.e., to the slow path or control plane) at the responder. In other words, the responder does not process the payload or 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]. In query messages, the label stack, as shown in Figure 3, carries both the forward and reverse path labels in the MPLS header. The GAL is still carried at the bottom of the label stack (with S=1).

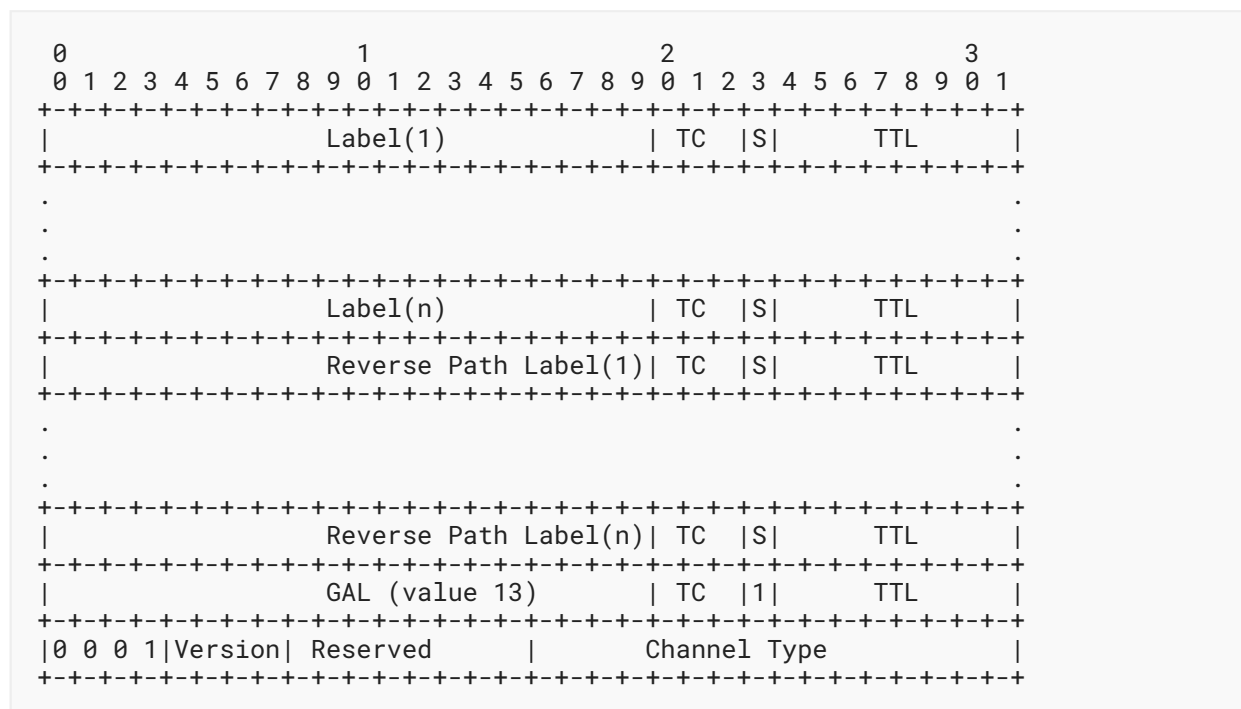


Figure 3: Example Query Message Header for an End-to-End SR-MPLS Policy in the Loopback Measurement Mode

5. Delay and Loss Measurement

5.1. Delay Measurement Message

As defined in [RFC6374], MPLS Delay Measurement (DM) query and response messages use the Associated Channel Header (ACH) (with the value 0x000C for delay measurement). This value identifies the message type and message payload that follow the ACH, as defined in Section 3.2 of [RFC6374]. For delay measurement, the same ACH value is used for both links and end-to-end SR-MPLS Policies.

5.2. Loss Measurement Message

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

- In the inferred mode, LM will measure the loss of specially generated test messages to infer the approximate data plane loss level. Inferred mode LM provides only approximate loss accounting.
- In the direct mode, LM will directly measure data plane packet loss. Direct mode LM provides perfect loss accounting but may require hardware support.

As defined in [RFC6374], MPLS LM query and response messages use the ACH (with the value 0x000A for direct loss measurement or value 0x000B for inferred loss measurement). This value identifies the message type and message payload that follow the ACH, as defined in Section 3.1 of [RFC6374]. For loss measurement, the same ACH value is used for both links and end-to-end SR-MPLS Policies.

5.3. Combined Loss/Delay Measurement Message

As defined in [RFC6374], combined LM/DM query and response messages use the ACH (with the value 0x000D for direct loss and delay measurement or the value 0x000E for inferred loss and delay measurement). The value identifies the message type and the message payload that follows the ACH, as defined in Section 3.3 of [RFC6374]. For combined loss and delay measurement, the same ACH value is used for both links and end-to-end SR-MPLS Policies.

5.4. Counters

The Path Segment Identifier (PSID) [RFC9545] **MUST** be carried in the received data packet for the traffic flow under measurement, in order to account for received traffic on the egress node of the SR-MPLS Policy. In the direct mode, the PSID in the received query message (as shown in Figure 4) can be used to associate the received traffic counter on the responder to detect the transmit packet loss for the end-to-end SR-MPLS Policy.

In the 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 the PSID for SR-MPLS Policy

The fields 0001, Version, Reserved, and Channel Type shown in Figure 4 are specified in [RFC5586].

Different values of the PSID can be used per Candidate-Path to account for received traffic and to measure packet loss at the Candidate-Path level. Similarly, different values of the PSID can be used per Segment List (SL) of the Candidate-Path for accounting received traffic to measure packet loss at the SL level. The same value of the PSID can be used for all SLs of the SR-MPLS Policy to measure packet loss at the SR-MPLS Policy level.

5.5. Block Number for Counters

The packet loss measurement using the Alternate-Marking Method defined in [RFC9341] may use the block number for data correlation for the traffic flow under measurement. As defined in Section 3.1 of [RFC9341], the block number is used to divide the traffic flow into consecutive blocks and count the number of packets transmitted and received in each block for loss measurement.

As described in Section 4.3 of [RFC9341], a protocol-based distributed solution can be used to exchange values of counters on the nodes for loss measurement. That solution is further described in this document using the LM messages defined in [RFC6374].

The querier node assigns a block number to the block of data packets of the traffic flow under measurement. The querier counts the number of packets transmitted in each block. The mechanism for the assignment of the block number is a local decision on the querier and is outside the scope of this document.

As an example, the querier can use the procedure defined in [RFC9714] for alternate marking of the data packets of the traffic flow under measurement. The responder counts the number of received packets in each block based on the marking in the received data packets. The querier

and responder maintain separate sets of transmit and receive counters for each marking. The marking can be used as a block number, or a separate block number can be incremented when the marking changes. Other methods can be defined for alternate marking of the data packets of the traffic flow under measurement to assign a block number for the counters.

The LM query and response messages defined in [RFC6374] are used to measure packet loss for the block of data packets transmitted with the previous marking, whereas data packets carry alternate marking. Specifically, LM query and response messages carry the transmit and receive counters (which are currently not incrementing) along with their block number to correlate for loss measurement.

Section 4.3 of [RFC9341] specifies that: "The assumption of this BN mechanism is that the measurement nodes are time synchronized." However, this is not necessary, as the block number on the responder can be synchronized based on the received LM query messages.

6. TLV Extensions

6.1. Return Path TLV Extension

In the two-way measurement mode, the responder may transmit the response message on a specific return path, for example, in an ECMP environment. The querier can request in the query message for the responder to send a response message back on a given return path (e.g., a co-routed bidirectional path). This allows the responder to avoid creating and maintaining additional states (containing return paths) for the sessions.

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

[RFC6374] defines query and response messages that can include one or more optional TLVs. This document defines the the Return Path TLV (5) to carry return path information in query messages. The Return Path TLV is specific to a measurement session. The format of the Return Path TLV is shown in Figure 5 below:

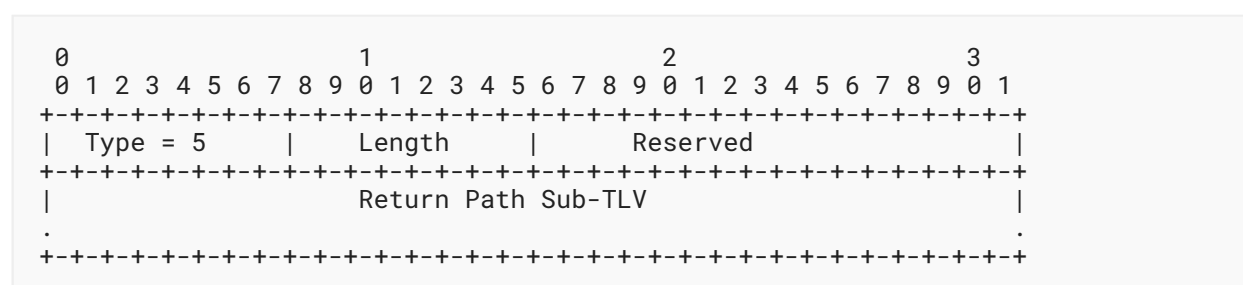


Figure 5: Return Path TLV

The Length is a one-byte field and is equal to the length of the Return Path Sub-TLV and the Reserved field in bytes. The Length **MUST NOT** be 0 or 1.

The Return Path TLV is defined in the "Mandatory TLV Type" registry space [RFC6374]. 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 the response message back on the return path specified in the Return Path TLV. The responder also **MUST NOT** add a Return Path TLV in the response message.

The Reserved field **MUST** be set to 0 and **MUST** be ignored on the receive side.

6.1.1. Return Path Sub-TLV Extension

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

- Type (value 1): MPLS Label Stack of the Return Path




Figure 6: MPLS Label Stack Sub-TLV in the Return Path TLV

The MPLS label stack contains a list of 32-bit LSEs that includes a 20-bit label value, an 8-bit TTL value, a 3-bit TC value, and a 1-bit End of Stack (S) field. An MPLS Label Stack Sub-TLV may carry a stack of labels or a Binding SID label [RFC8402] of the Return SR-MPLS Policy.

The Length is a one-byte field and is equal to the length of the label stack field and the Reserved field in bytes. The Length **MUST NOT** be 0 or 1.

The Return Path TLV **MUST** carry only one Return Path Sub-TLV. The MPLS Label Stack in the Return Path Sub-TLV **MUST** contain at least one MPLS Label. 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 **MUST** send the 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.

6.2. Block Number TLV Extension

[RFC6374] defines query and response messages that can include one or more optional TLVs. This document defines the Block Number TLV (6) to carry (8-bit) Block Number 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 Length is a one-byte field and is equal to 2 bytes.

The Block Number TLV is defined in the "Mandatory TLV Type" registry space [RFC6374]. 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 Block Number TLV is specific to a measurement session. The R flag is used to indicate the query and response message direction associated with the Block Number. The R flag **MUST** be clear in the query message for the Block Number associated with Counter 1 and Counter 2, and set in the response message for the Block Number associated with Counter 3 and Counter 4.

The Reserved field **MUST** be set to 0 and **MUST** be ignored on the receive side.

7. ECMP for SR-MPLS Policies

The SLs of an SR-MPLS Policy can have ECMPs between the source and transit nodes, between transit nodes, and between transit and destination nodes. Usage of a node-SID [RFC8402] by the SLs of an SR-MPLS Policy can result in ECMP paths. In addition, usage of an Anycast-SID [RFC8402] by the SLs of an SR-MPLS Policy can result in ECMP paths via transit nodes that are part of that anycast group. The query and response messages are sent to traverse different ECMP paths to measure the delay of each ECMP path of an SL.

The forwarding plane has various hashing functions available to forward packets on specific ECMP paths. For end-to-end SR-MPLS Policy delay measurement, different entropy label values [RFC6790] can be used in query and response messages to take advantage of the hashing function in the 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.

8. Extended TE Link Metrics Advertisement

The extended TE metrics for link delay (namely, average delay, minimum delay, maximum delay, and delay-variance) and packet loss can be computed using the performance measurement procedures described in this document and can be advertised in the routing domain as follows:

- For OSPF, IS-IS, and the Border Gateway Protocol - Link State (BGP-LS), the 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 [RFC9356] and IS-IS [RFC8668] using the same protocol extensions defined in [RFC7471] and [RFC8570], respectively.
- The advertised delay-variance metric is computed as Packet Delay Variation (PDV), as described in Section 4.2 of [RFC5481].

9. Backwards Compatibility

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

10. Manageability Considerations

The manageability considerations described in Section 7 of [RFC6374] and Section 6 of [RFC7876] are applicable to this specification.

11. Security Considerations

The security considerations specified in [RFC6374], [RFC7471], [RFC8570], [RFC8571], [RFC7876], and [RFC9341] also apply to the procedures described in this document.

The procedure defined in this document is intended for deployment in a single operator administrative domain. As such, the querier node, the responder node, the forward path, and the return paths are provisioned by the operator for the probe session. It is assumed that the operator has verified the integrity of the forward and return paths of the probe packets.

The Return Path TLV extensions defined in this document may be used for potential address spoofing. For example, a query message may carry a return path that has a destination that is not local at the querier. To prevent such possible attacks, the responder may drop the query messages when it cannot determine whether the return path has the destination local at the

querier. The querier may send a proper source address in the Source Address TLV. The responder can use this to make that determination, for example, by checking the access control list provisioned by the operator.

12. IANA Considerations

IANA has allocated values for the following Mandatory TLV Types [RFC6374] from the "MPLS Loss/Delay Measurement TLV Object" registry contained within the "Generic Associated Channel (G-ACh) Parameters" registry group:

Code	Description	Reference
5	Return Path	RFC 9779
6	Block Number	RFC 9779

Table 1: MPLS Loss/Delay Measurement TLV Types

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

IANA has created the "Return Path Sub-TLV Types" registry. All code points are allocated per the registration policies shown in Table 2 (see [RFC8126]).

Value	Description	Reference
1 - 175	IETF Review	RFC 9779
176 - 239	First Come First Served	RFC 9779
240 - 251	Experimental Use	RFC 9779
252 - 254	Private Use	RFC 9779

Table 2: Return Path Sub-TLV Type Registry

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

Value	Description	Reference
0	Reserved	RFC 9779
1	MPLS Label Stack of the Return Path	RFC 9779
255	Reserved	RFC 9779

Table 3: Return Path Sub-TLV Types

13. References

13.1. 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>>.
- [RFC9341] Fioccola, G., Ed., Cociglio, M., Mirsky, G., Mizrahi, T., and T. Zhou, "Alternate-Marking Method", RFC 9341, DOI 10.17487/RFC9341, December 2022, <<https://www.rfc-editor.org/info/rfc9341>>.

13.2. Informative References

- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <<https://www.rfc-editor.org/info/rfc3031>>.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001, <<https://www.rfc-editor.org/info/rfc3032>>.
- [RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., Ed., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, DOI 10.17487/RFC5082, October 2007, <<https://www.rfc-editor.org/info/rfc5082>>.
- [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>>.
- [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", RFC 5586, DOI 10.17487/RFC5586, June 2009, <<https://www.rfc-editor.org/info/rfc5586>>.

-
- [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>>.
- [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>>.
- [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>>.
- [RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.
- [RFC9356] Talaulikar, K., Ed. and P. Psenak, "Advertising Layer 2 Bundle Member Link Attributes in OSPF", RFC 9356, DOI 10.17487/RFC9356, January 2023, <<https://www.rfc-editor.org/info/rfc9356>>.
- [RFC9545] Cheng, W., Ed., Li, H., Li, C., Ed., Gandhi, R., and R. Zigler, "Path Segment Identifier in MPLS-Based Segment Routing Networks", RFC 9545, DOI 10.17487/RFC9545, February 2024, <<https://www.rfc-editor.org/info/rfc9545>>.
- [RFC9714] Cheng, W., Ed., Min, X., Ed., Zhou, T., Dai, J., and Y. Peleg, "Encapsulation for MPLS Performance Measurement with the Alternate-Marking Method", RFC 9714, DOI 10.17487/RFC9714, February 2025, <<https://www.rfc-editor.org/info/rfc9714>>.
- [IEEE802.1AX] IEEE, "IEEE Standard for Local and Metropolitan Area Networks - Link Aggregation", IEEE Std 802.1AX-2020, DOI 10.1109/IEEESTD.2020.9105034, May 2020, <<https://doi.org/10.1109/IEEESTD.2020.9105034>>.
-

Acknowledgments

The authors would like to thank Thierry Couture and Ianik Semco for the discussions on the use cases for performance measurement in segment routing networks. The authors would like to thank Patrick Khordoc, Ruby Lin, and Haowei Shi for implementing the mechanisms defined in this document. The authors would like to thank Greg Mirsky and Xiao Min 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 the MPLS expert review; Zhaohui Zhang for the RTGDIR early review; Tony Li for shepherd's review; Ned Smith for the SECDIR review; Roni Even for the Gen-ART review; Marcus Ihlar for the TSV-ART review; Dhruv Dhody for the OPSDIR review; and Gunter Van de Velde, Paul Wouters, Éric Vyncke, Murray Kucherawy, John Scudder, and Roman Danyliw for the IESG 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