

Benchmarking Methodology for MPLS Segment Routing

draft-vfv-bmwg-srmpls-bench-meth-01

Hybrid, Mar 2022, IETF 113

Giuseppe Fioccola (Huawei)

Eduard Vasilenko (Huawei)

Paolo Volpato (Huawei)

Luis Contreras (Telefonica)

Background

RFC5695 describes a methodology specific to the benchmarking of MPLS forwarding devices.

Segment Routing (SR), defined in RFC8402, leverages the source routing paradigm and can be applied to MPLS data plane (**SR-MPLS**)

- However, there is no standard method defined to compare and contrast the foundational SR-MPLS packet forwarding capabilities of network devices.
- This new I-D aims to complement RFC5695 by defining the methodology for benchmarking SR-MPLS.
- It builds upon RFC2544, RFC5695 and RFC8402.

SR-MPLS Operations

SR can be directly applied to the MPLS architecture with no change to the forwarding plane (RFC8660).

- A segment is encoded as an MPLS label.
- An SR Policy is instantiated as a stack of labels.

PUSH, NEXT, and CONTINUE are the basic SR-MPLS operations applied by the forwarding plane.

- PUSH consists of the insertion of a segment at the top of the segment list. The top of the segment list is the outer label of the label stack.
- NEXT consists of the inspection of the next segment. The active segment is completed and the next segment becomes active.
- CONTINUE happens when the active segment is not completed; hence, it remains active.

Test Methodology (1/2)

Similarity with RFC5695:

- The Device Under Test (DUT) is connected to the test ports on the test tool according to RFC2544.
- The recommended topology for SR-MPLS Forwarding Benchmarking should be the same as MPLS and it is described in RFC5695 for both single-port and multi-port scenarios.
- The tests for SR-MPLS will use the Frame characteristics as described in RFC5695.

Test Methodology (2/2)

Initial proposed changes from RFC5695:

- It is RECOMMENDED that all of the ports on the DUT and test tool support a SR extension for dynamic IGP such as IS-IS (RFC8667) and OSPF (RFC8665) as well as BGP (RFC8669).
- RFC5695 requires exactly a single entry in the MPLS label stack in an MPLS packet. With SR-MPLS the number of entries in the label stack SHOULD be generalized in order to benchmark SR Policy.
 - But the number of labels MUST be chosen taking into account the effective maximum frame payload size (payload of 1500 octets for Ethernet).
- There are new parameters that MUST be added to the parameters specified in RFC5695:
 - SR-MPLS Forwarding Operations (PUSH/ NEXT/ CONTINUE),
 - Number of Segments considered in the MPLS Label Stack,
 - Global SIDs or Local SID forwarding behavior,
 - SR Policy headend or endpoint behavior

SR-MPLS Forwarding Benchmarking Tests

RFC5695 defines tests (Throughput, Latency, Frame Loss, System Recovery, Reset) for **MPLS Label Push**, **MPLS Label Swap** and **MPLS Label POP**.

Similarly, for **SR-MPLS** three fundamental operations needs to be tested: **PUSH**, **NEXT** and **CONTINUE**

- The PUSH operation corresponds to the Label Push function. It consists of pushing one or more MPLS labels on top of an incoming packet then sending it towards a particular next hop.
- The NEXT operation corresponds to the Label Pop function, that consists of removing the topmost label.
- The CONTINUE operation corresponds to the Label Swap function. It consists of forwarding the packet on the outgoing interface based on the incoming label.

Therefore the procedure can be similar to RFC5695 but with a different Reporting Format

Next Steps

Feedback is welcome on some pending points:

- Our current approach is to extend RFC5695 and RFC2544. The alternative is to rewrite more text from RFC5695 and RFC2544 to create one more convenient document for the reader
- RFC5695 includes only 1 label in the stack. But it may be needed to test more labels to support SR policy with reasonable depth (number of labels MUST be reported)
- The tests are under the condition of not having additional background traffic in place. But it could be valuable to add a test scenario including complementary background traffic.
- Should we consider the TE test out of scope for this document? TE was typically in separate documents in BMWG (e.g. RFC5695 and RFC6894).

Welcome inputs, comments

Open to new coauthors, contributors

Thank you