

Benchmarking Methodology for MPLS Segment Routing

draft-vfv-bmwg-srmpls-bench-meth-06

IETF 117 Update

Luis Contreras (Telefonica)

Bruno Decraene (Orange)

Giuseppe Fioccola (Huawei)

Eduard Vasilenko (Huawei)

Paolo Volpato (Huawei)

IETF 117, San Francisco, July 26

Recap of Draft's Target

- There is no standard method to benchmark the foundational Segment Routing (SR) [RFC8402] packet forwarding capabilities of network devices.
- SR leverages the source routing paradigm.
 - The headend node steers a packet through an SR Policy [I-D.ietf-spring-segment-routing-policy], instantiated as an ordered list of segments.
 - A segment is referred to by its Segment Identifier (SID).
 - SR supports per-flow explicit routing while maintaining per-flow state only at the ingress nodes to the SR domain.
- The SR architecture can be instantiated on two data-plane:
 - SR over MPLS (SR-MPLS), and
 - SR over IPv6 (SRv6).
- This document is limited to SR-MPLS. It aims to extend the efforts of [RFC1242], [RFC2544] and [RFC5695] to an SR-MPLS network.

SR-MPLS Forwarding Benchmarking Tests

- An SR Policy is instantiated through the MPLS Label Stack: the Segment IDs (SIDs) of a Segment List are inserted as MPLS Labels.
- The forwarding functions available for MPLS networks allow implementing the SR operations. SR-MPLS applies three operations on the forwarding plane:
 - PUSH [Label Push]. One or more MPLS labels are pushed on top of an incoming packet, before the packet is sent out of a physical/virtual interface.
 - NEXT [Label Pop]. The topmost label is removed. The next action depends on the instruction associated with the active SID. It equals to Penultimate Hop Popping (PHP).
 - CONTINUE [Label Swap]. It associates an incoming label with an outgoing interface and outgoing label. The packet is forwarded to the outgoing interface. It is equivalent to Ultimate Hop Popping (UHP).
- The benchmark procedure can be similar to RFC5695 with extensions:
 - Test SID list longer than 1 SID (2 are recommended, many are optional).
 - Different Reporting Format.
 - At least one protocol for the SID population is recommended (ISIS or OSPF or SR Policy).
- The tests (throughput, buffer size, latency, etc.) are repeated for every operation.

Draft's History (main changes)

- Version -00 submitted on March 2022
- Version -01 presented at IETF 113 – Initial test methodology discussed
- Version -02 presented at IETF 114 – Incorporated comments from the chairs and the list to review the test setup and methodology, trial duration and traffic verification
- Versions -03 and -04 submitted in October 2022 and presented at IETF 115 – Included further comments on buffer's size test, ways for equalizing the link's load balancing
- Version -05 submitted on February 2023 – Addressed comments on allowing a longer list of SID as an optional test. Bruno joined as coauthor.
- Version -06 uploaded on March 2023 and presented at IETF 116.
- Since then draft stable.

Next Steps

- As requested at IETF 116, we think the draft is stable enough for moving to WG adoption.
- At IETF 116 a final round of reviews was proposed to confirm it.
- People who provided valuable comments to the initial versions of the draft may acknowledge that the comments have been addressed.
- A WG adopted document would help the draft to get more attention from companies doing tests.
- 3rd party tests are also desired before it would become RFC.

Thank you