

Benchmarking Methodology for Segment Routing MPLS

draft-vfv-bmwg-srmpls-bench-meth-04

Hybrid, Nov 2022, IETF 115

Giuseppe Fioccola (Huawei)

Eduard Vasilenko (Huawei)

Paolo Volpato (Huawei)

Luis Contreras (Telefonica)

Background

- **Segment Routing (SR)**, defined in **RFC 8402**, leverages the source routing paradigm.
 - The headend node steers a packet through an SR Policy instantiated as an ordered list of segments.
 - **SR-MPLS**, focus of this draft, does not require any change to the MPLS forwarding plane. A Prefix-SID is allocated in the form of an MPLS label.
- **RFC 5695** describes a methodology for benchmarking MPLS forwarding devices.
 - It considers the most common MPLS packet forwarding scenarios and corresponding performance measurements.
- The purpose of this draft is to describe a methodology specific to the benchmarking of **SR-MPLS** capable devices.
 - The methodology described is a complement for RFC 5695.
 - It builds upon RFC 2544, RFC 8402, and RFC 5695.
 - RFC 4814, RFC 8219, and RFC 9004 are also key references.

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 – Corresponds to the Label Push function. 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 – Corresponds to the Label Pop function. The topmost label is removed. The action before and/or after popping depends on the instruction associated with the active SID. It equals to Penultimate Hop Popping (PHP).
 - CONTINUE – Corresponds to the Label Swap function. 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 overall procedure can be similar to RFC5695 with some extensions:

- Test SID list longer than 1 SID (2 are recommended, many are possible)
- Different Reporting Format
- At least one protocol for the SID population is recommended (ISIS or OSPF or SR Policy).

Changes from -02

Two versions (-03 and -04) submitted in October 2022, to address the comments received at IETF 114 and on the list from Gabor Lencse and Boris Khasanov:

- Buffer's size test (back-to-back) is added on the basement of RFC 9004
- Reference to RFC 4814 on how to deal with L2 links that may have staffing (no need for the default Ethernet link)
- Reference to RFC 4814 on how to deal with address randomization for equalization of link's load-balancing (no need for SR itself but needed for IP addresses and transport layer ports)
- Reference to RFC 8219 on how to improve latency measurement if needed (optional)
- Better alignment with the rules on how to use reserved IETF words
- Stronger statement that special capabilities just for tests are prohibited
- SR policies may be installed by many protocols
- Make at least routing protocol mandatory
- Increase for the test time
- Improvement for Reset test methodology (RFC 5695)
- Editor changes

Next Steps

BMWG may consider adopting this draft

Welcome inputs, comments

Open to new coauthors, contributors

Thank you