

Benchmarking Methodology for MPLS Segment Routing

draft-vfv-bmwg-srmppls-bench-meth-08

IETF 118 Update

Luis Contreras (Telefonica)

Bruno Decraene (Orange)

Giuseppe Fioccola (Huawei)

Eduard Vasilenko (Huawei)

Paolo Volpato (Huawei)

IETF 118, Prague, November 6

Draft Overview

Problem:

- Segment Routing (**SR**) (RFC 8402) leverages the source routing paradigm and can be applied to both the MPLS data plane (**SR-MPLS**) and the IPv6 data plane (**SRv6**).
- However, there is **no standard method** defined to compare **SR packet forwarding capabilities** of network devices.

Solution:

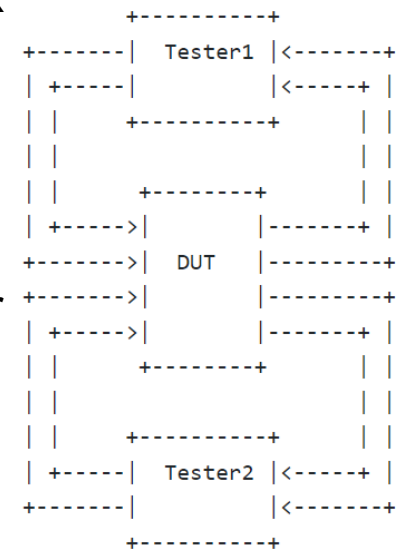
- This document defines a **methodology for benchmarking SR-MPLS**.
- It complements RFC 5695.
- It builds upon RFC 2544, RFC 5695, and RFC 8402. RFC 4814, RFC 8219, and RFC 9004 are included for refinement of some specific points.

Update from IETF 117

- Version -06 (March 2023): discussed at IETF 116 and IETF 117.
- Version -07 (September 2023): keep-alive edition.
- Version -08 (October 2023):
 - Extensively reviewed to incorporate almost all relevant text from referenced RFCs in order to be a stand-alone document to read (except buffering tests that is too long to copy from RFC 9004).
 - This should address a comment from Carsten at the mic.
 - Removed duplicated text about the technicalities of SR-MPLS in section 2.
 - Clarified that the MPLS label stack in scope of this document has a minimum of two entries (2 SIDs), but it is RECOMMENDED that the tests are applied to label stacks with more than two SIDs.
 - As indicated by Bruno, some tests have to be repeated with N labels/SIDs. This is necessary to stress the DUT and understand the maximum SID list supported.
 - Expanded section 4 (Reporting Format) to list all the parameters to be recorded in the tests.
 - Fixed nits and errors.

SR-MPLS Forwarding Benchmarking Tests

- The performance of a modern packet forwarding engines may be so high that 2 or more testers are needed to sufficiently load the DUT.
- RFC 5695, sec. 6, specifies the test tool(s) send traffic from the Tx interface(s) to the DUT at constant load for a fixed-time interval.
- If any frame loss is detected, a new iteration is started with a decreased load to determine the maximum offered frame rate with a zero frame loss (No-Drop Rate - NDR).
- The other parameters (test duration, number of interfaces, number of addresses, frame size, etc.) remain constant.
- The test can be repeated with a varying number of Segments pushed on ingress in order to measure the resulting maximum number. It can also be tested the maximum number of Segments that are correctly load-balanced.
- Therefore, the two main parameters that can be evaluated are:
 - Maximum offered frame rate,
 - Maximum number of Segments that can be pushed and hashed by the SR node for load-balancing.



Next Steps

- As already requested at IETF 117, the authors believe the draft is stable enough for WG adoption.
- Informally, some replies in the mail alias suggested we can move ahead with it.
- As already discussed, the WG adoption would help the draft to get more attention from companies to run validation tests.
- 3rd party tests are also desired before it would become RFC.

Thank you