Benchmarking Methodology for MPLS Segment Routing

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

IETF 116, Hybrid, March 203

Luis Contreras (Telefonica)
Bruno Decraene (Orange)
Giuseppe Fioccola (Huawei)
Eduard Vasilenko (Huawei)
Paolo Volpato (Huawei)

Recap of Draft's Target

- There is no standard method to compare the foundational SR packet forwarding capabilities of network devices.
- Segment Routing [RFC8402], 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.
- This document aims to extend the efforts of [RFC1242], [RFC2544] and [RFC5695] to SR network.
- 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.

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 some 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

- 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:
  - Revised Test Setup and Methodology
  - Added new sections on Protocol Addresses, Trial Duration and Traffic Verification
  - Included additional consideration on the relationship with RFC5695 and RFC2544
  - Traffic Engineering and Services (VPNs) have been put out of the scope
  - References to RFC9004 and ETSI GR NFV TST 007.
- Versions -03 and -04 submitted in October 2022, -04 presented at IETF 115 – Included further comments from Gabor and Boris:
  - Buffer's size test
  - Reference to RFC 4814 for L2 links with staffing and address randomization for equalization of link's load balancing
  - Reference to RFC 8219 on how to improve latency measurement, and others.
- Version -05 submitted on February 2023 – Included comments from Bruno. Bruno also joined as coauthor.
  - Longer list of SID as an optional test
  - Editorial changes.
- Version -06 uploaded on March to address few comments from Gabor and to align with SRv6 draft.

Next Steps

- The document looks stable.

We are asking for BMWG draft adoption.

- 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