Benchmarking Methodology for IPv6 Segment Routing

draft-vfv-bmwg-srv6-bench-meth-01

Hybrid, Mar 2022, IETF 113

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

Background

RFC5180 defines the methodology for the benchmarking of IPv6 forwarding devices.

Segment Routing (SR), defined in RFC8402, leverages the source routing paradigm and can be applied to IPv6 data plane (**SRv6**)

- However, there is no standard method defined to compare and contrast the foundational SRv6 packet forwarding capabilities of network devices

- This new I-D aims to complement RFC5180 by defining a methodology for benchmarking SRv6.

- It builds upon RFC2544, RFC5180 and RFC8402.

SRv6 Operations

SR can be applied to the IPv6 architecture:

- A new type of routing header called SRH has been defined in RFC8754.
- An segment is encoded as an IPv6 address (SRv6 SID).
- An SR Policy is instantiated as an ordered list of SRv6 SIDs in the SRH.

There are three different categories of nodes that may be involved in SRv6 networks:

- The SR source node is the headend node and steers a packet into an SR Policy. It sets the first SID of the SR Policy as IPv6 Destination Address of the packet.

- The SR transit node forwards packets destined to a remote segment as a normal IPv6 packet on the basis of the IPv6 destination address, because the IPv6 destination address does not locally match with a segment.

- The SR segment endpoint node receives packets whose IPv6 destination address is locally configured as a segment. For each SR packet, it inspects the SRH and replaces the IPv6 destination address with the new active segment.

SRv6 Network Programming

In addition to the basic SRv6 packet processing, the SRv6 Network Programming model (RFC8986) describes a set of functions that can be associated to segments and executed in a given SRv6 node.

According to RFC8986, 128 bit SID can be logically split into three fields:

- LOCATOR:FUNCTION:ARGS, where 128=L+F+A
  - The LOC corresponds to an IPv6 prefix (for example with a length of 48, 56 or 64 bits) that can be distributed by the routing protocols.
  - The FUNCT code identifies a function residing in a node
  - The ARG bits are used to provide information (arguments) to a function.

A list of Behaviors is specified in RFC8986 and a subset of this function could be tested to benchmark the SRv6 Network Programming capability

Test Methodology (1/2)

Similarity with RFC5180:

- The Device Under Test (DUT) is connected to the test ports on the test tool according to RFC2544.
- The test topology recommended for the SRv6 performance evaluation are the same as IPv6 and are described in RFC5180 and RFC2544, in both single-port and multi-port scenarios.
- The tests for SRv6 will use the Frame characteristics as described in RFC5180.

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 Segment Routing extension for dynamic IGP for routing such as IS-IS and OSPF as well as BGP.

- RFC5180 recommends a limited extension header (EH) chain for testing: Routing header (24-32 bytes) + Destination options header (8 bytes) + Fragment header (8 bytes). This is not enough and does not fit well for SRv6:
  - The length of the SRH is [n x 16 + 8] bytes, where n is the number of segments.
  - n should be > 1 in order to benchmark SR Policy
  - For SRv6, the EH chain length that is used MUST be reported and the DUT MUST traverse the chain of EHs, so the potential impact on performance can be observed.

- There are new parameters that MUST be added to the parameters specified in RFC5180 and RFC2544:
  - SRv6 types of nodes,
  - Number of Segments considered in the SRH,
  - Extension header chain (including SRH) characteristics and length,
  - Global SIDs or Local SID forwarding behavior,
  - SR Headend or Endpoint Behaviors eventually associated with a SID, as specified in RFC8986.

SRv6 Forwarding Benchmarking Tests

RFC5180 defines tests (Throughput, Latency, Frame Loss, System Recovery, Reset) for basic **IPv6** forwarding

For **SRv6**, new tests need to be added to characterize a **SRv6 Source Node**, a **SRv6 Segment Endpoint Node**, a **Transit Node**

- The processing of the SR source node corresponds to the insertion of the SRH, with SIDs stored in reverse order, and setting of the IPv6 DA as first SID of the SR Policy.
- The processing of the SR segment endpoint node corresponds to the detection of the new active segment, modification of the IPv6 DA of the IPv6 header, and forwarding of the packets.
- The processing of the SR transit node corresponds to normal forwarding of the packets containing the SR header. In SRv6 the transit nodes do not need to be SRv6 aware.

Therefore new tests are defined, the overall procedure can be similar to RFC5180 but with a different Reporting Format

Next Steps

Feedback is welcome on some pending points:

- Our current approach is to extend RFC5180 and RFC2544. The alternative is to rewrite more text from RFC5180 and RFC2544 to create one more convenient document for the reader

- What “Functions” should be included? Should we consider only End, End.X, End.T or those defined in RFC8986? In any case the tested Behaviors MUST be reported.

- Should we cover CSID (Compressed SID) tests? It could be done later considering that the corresponding I-D (draft-ietf-spring-srv6-srh-compression) is still in progress.

- 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