

Poster: A First-Look at Segment Routing Deployment in a Large European ISP

Emeline Marechal[†], Yining Shao^{*}, Marc Bruyère[‡], Benoit Donnet[†]
[†] Université de Liège (Belgium), ^{*}Waseda University (Japan), [‡] IIJ (Japan)

Context

In a nutshell, Segment Routing [4] (SR) is a loose source routing paradigm based on an ordered list of *segments* (i.e., one or more forwarding instructions). Over the years, SR has found a suitable usage in network monitoring, traffic engineering, or failure recovery [6], among others. Two forwarding plane encapsulations are proposed for SR: MPLS (the focus of this extended abstract – SR-MPLS) and Extension Headers for IPv6 (SRv6). SRv6 deployment in the Internet has been investigated by [7]. To the best of our knowledge, this extended abstract is the first attempt in revealing the deployment and usage of SR-MPLS in the wild. Early results suggest that SR-MPLS is mainly used in interworking with classic MPLS tunnels in Vodafone.

Background

SR defines multiple types of segments, but the two most common are *node segments* and *adjacency segments*. A *node segment* (see Fig. 1) represents the IGP least cost path between any router and a specified prefix. These segments can contain one or multiple IGP hops and have domain-wide significance. An *adjacency segment* represents an IGP adjacency between two routers and will cause a packet to traverse that specified link. These segments only have local significance.

In SR-MPLS, each segment is identified by a unique number, a *Segment Identifier* (SID) implemented as an MPLS 20-bit label value. For *node segments*, labels are globally allocated in the ISP domain from the *SR Global Block* (SRGB), a range of MPLS labels used solely for SR. Within the domain, each SR-capable router may reserve a different range for its SRGB, although it is advised to use the same on every router for simplicity. For instance router A might reserve labels 16,000 to 17,000 while router B reserves labels 20,000 to 21,000. Both routers map SIDs to MPLS labels by adding the SID to the lowest SRGB value. Therefore, router A maps (for instance) SID 1 to MPLS label 16,001, while router B maps the same SID to MPLS label 20,001. For *adjacency segment*, the MPLS labels are allocated automatically from the dynamic label range, without concern for domain-wide coordination.

When a SR Ingress router receives a packet, it encapsulates it into an MPLS stack (R1 on Fig. 1). Each label within the

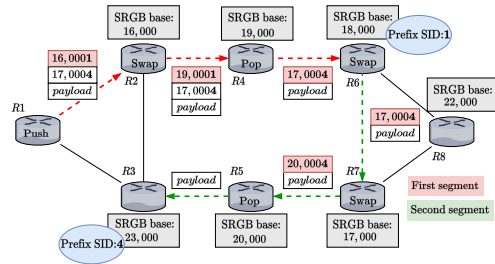


Figure 1: Node Segment example. R1 pushes into the packet a label stack with two labels: the outer one representing the node segment to R6, and the inner one representing the node segment to R3. R1 computes the inner label by adding R3’s SID (4) to R7’s SRGB base (17,000), and it computes the outer label by adding R6’s SID (1) to R2’s SRGB base (16,000).

stack represents a particular segment (made of one or more IGP hops). Following usual MPLS operations, the stack top label is used, within a segment, to forward packets. When the packet reaches the end of the segment, the router pops the top of the stack and forwards the packet to the next segment, and so on, until the packet leaves the SR domain.

In this extended abstract, we present preliminary measurement results of the SR-MPLS deployment (focusing on node segment) in a large European ISP, Vodafone Germany (ASN 3209), as it is known for having deployed SR-MPLS [2]. We rely on a longitudinal dataset collected by CAIDA with TNT [5], a Paris traceroute extension able to reveal all MPLS tunnels. TNT also provides a fingerprint for each collected IPv4 address, which allows us to associate the hardware vendor to the device [8].

Methodology

We identify SR in the TNT dataset using the MPLS label stack quoted in the ICMP time-exceeded message [5]. We can only see labels for *explicit* tunnels, which are fully visible with traceroute. Other types of tunnels exist, namely, *opaque*, *invisible*, and *implicit* tunnels [5] that cannot be spotted with traceroute. Even though traceroute does not see such tunnels, TNT can infer their presence and reveal them, except from their labels, meaning we cannot infer if they use SR or not.

This work is supported by the CyberExcellence project funded by the Walloon Region, under number 2110186.

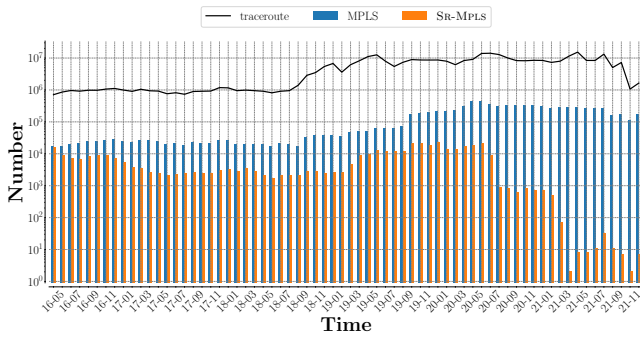


Figure 2: SR-MPLS deployment in Vodafone. The number of traceroutes per month lies between 10^6 and 10^7 . Our method observed more MPLS than SR-MPLS tunnels, especially after mid-2020, where the number of SR-MPLS tunnels dropped by 70%.

Recommendation by Cisco is to use the label range 16,000 to 23,999 for the SRGB. Additionally, the range 15,000 to 15,999 is also reserved for manual allocations of labels on Cisco devices [1]. These are default values implemented in their hardware, and operators should have little motivation to change it, if they want to avoid making their SR deployment any more complex. Therefore, we base our preliminary identification of SR on the observation of this range of values in MPLS labels.

Given that our visibility is limited to *node segments* for *explicit* tunnels on Cisco hardware, our results represent a lower bound of SR deployment.

Preliminary results

CAIDA collects TNT data towards all routed /24 prefixes ($\sim 10M$). We consider data between May 2016 and December 2021. Fig. 2 presents an overview of SR-MPLS deployment in Vodafone. The number of traceroutes per snapshot lies between 10^6 and 10^7 . Our method observed more MPLS than SR-MPLS tunnels, especially after mid-2020, where the number of SR-MPLS tunnels dropped by 70%. We can only speculate on the reason for this drop, but it may be due to Vodafone transitioning to use *invisible* tunnels.

We further investigated if the hardware seen for the SR hops does correspond to Cisco, which would be consistent with our assumptions. Fig. 3 shows the hardware distribution for the different types of traffic (SR, MPLS, IPv4, and globally) in Vodafone. In 80% of the cases, IP addresses corresponding to SR-MPLS also correspond to Cisco devices, reinforcing our assumption for identifying SR.

Next, we explore SR-MPLS tunnel lengths. We already know that classic MPLS tunnels tend to be quite short [9]. Surprisingly, 90% of traces presented only a single SR hop along the path. 93% of the time, this SR hop actually belonged to a longer MPLS tunnel. We typically saw it as the last hop

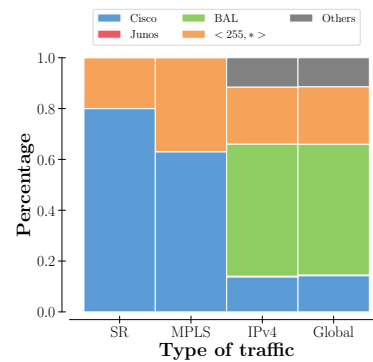


Figure 3: Hardware distribution for each type of traffic. The interfaces can be classified into Cisco, Junos, BAL (Brocade, Alcatel, Linux), the incomplete signature $< 255, * >$, or Others. 80% of interfaces seen in SR traffic belongs to a Cisco device. The remaining 20% present the incomplete signature $< 255, * >$, which most probably corresponds to Cisco devices that did not answer the second probe necessary for identification.

of the tunnel, but the first position, or, rarely, the middle of the tunnel, can also be possible. These observations might suggest that SR-MPLS is currently used in interworking with a classic MPLS deployment, for example with a *Mapping Server* [3]. It is quite reasonable to assume that SR needs incremental deployment, as replacing hardware is expensive.

Our next steps are to extend our study to other ASes and seek validation of our inferences from network operators.

REFERENCES

- [1] Cisco. 2021. *Segment Routing Configuration Guide for Cisco ASR 9000 Series Routers*. Cisco Press, Chapter Chap. 5.
- [2] J. Davidson. [n. d.]. Segment Routing Blazes a Trail in Germany! ([n. d.]). <https://blogs.cisco.com/sp/segment-routing-blazes-a-trail-in-germany> Last Access: August, 17th 2022.
- [3] A. Bashandy et al. 2019. *Segment Routing MPLS Interworking with LDP*. RFC 8661. Internet Engineering Task Force.
- [4] C. Filsfil et al. 2018. *Segment Routing Architecture*. RFC 8402. Internet Engineering Task Force.
- [5] J. R. Luttringer et al. 2020. Let There Be Light: Revealing Hidden MPLS Tunnels with TNT. *IEEE Transactions on Network and Service Management (TNSM)* 17, 2 (June 2020), 1239–1253.
- [6] P. L. Ventre et al. 2020. Segment Routing: A Comprehensive Survey of Research Activities, Standardization Efforts, and Implementation Results. *IEEE Communications Surveys & Tutorials* 23, 1 (November 2020), 182–221.
- [7] V.-A. Padurean et al. 2022. SRv6: Is There Anybody Out There?. In *Proc. International Workshop on Traffic Measurements for Cybersecurity (WTMC)*.
- [8] Y. Vanaubel et al. 2013. Network Fingerprinting: TTL-Based Router Signature. In *Proc. ACM Internet Measurement Conference (IMC)*.
- [9] Y. Vanaubel et al. 2015. MPLS Under the Microscope: Revealing Actual Transit Path Diversity. In *Proc. ACM Internet Measurement Conference (IMC)*.