Autonomous Systems under AReST: Advanced Revelation of Segment Routing Tunnels

Florian Dekinder
University of Liège
Liège, Belgium
florian.dekinder@uliege.be

Kevin Vermeulen
LIX, CNRS, École Polytechnique
France

Benoit Donnet
University of Liège
Liège, Belgium
benoit.donnet@uliege.be

Abstract

Segment Routing (SR), an advanced source routing mechanism, is a promising technology with a wide range of applications that has already gained traction from hardware vendors, network operators, and researchers alike. However, despite the abundance of activity surrounding SR, little is known about how to gauge SR deployment and its usage by operators.

This paper introduces a methodology, called **A**dvanced **Re**velation of **S**egment Routing **T**unnels (AReST), for revealing the presence of SR with MPLS as forwarding plane (SR-MPLS). AReST relies on standard measurement tools, like `traceroute` and fingerprinting, and post-processes the collected data for highlighting evidence of SR-MPLS. Our results show that AReST is efficient in revealing the presence of SR-MPLS in various autonomous systems, achieving perfect precision on our ground truth directly obtained from an operator. We also make a preliminary characterization of the SR-MPLS deployment and show that it is commonly deployed within Content, Transit, and Tier-1 providers and, occasionally, in interworking with classic MPLS. The data collected, as well as our source code, are available to the research community.

CCS Concepts

• **Networks** → Network layer protocols; **Network measurement**; **Topology analysis and generation**.

Keywords

Segment Routing, MPLS, traceroute, fingerprinting, labels, AReST

ACM Reference Format:
Florian Dekinder, Kevin Vermeulen, and Benoit Donnet. 2025. Autonomous Systems under AReST: Advanced Revelation of Segment Routing Tunnels. In *Proceedings of the 2025 ACM Internet Measurement Conference (IMC '25), October 28–31, 2025, Madison, WI, USA.* ACM, New York, NY, USA, 17 pages. https://doi.org/10.1145/3730567.3764436

1 Introduction

Segment Routing [12, 29, 30, 32] (SR) is a technology based on the *source routing* paradigm, where a source node determines the path that a packet takes through the network by encoding a sequence of instructions, called *segments*, into the packet header. SR is not limited to a sequence of forwarding instructions, but can also chain several services to obtain complex behaviors. Each segment can either enforce a topological requirement (e.g., go through a node or an interface) or a service requirement (e.g., execute an operation on the packet).

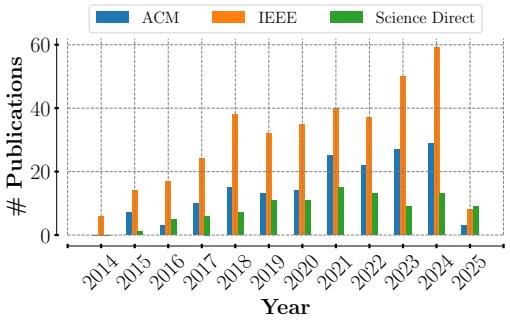

Figure 1: Number of publications involving Segment Routing between 2014 and 2025. Data were collected on March 31st, 2025, with "Segment Routing" as keyword on several online scientific libraries.

Although its first standardization occurred in 2018 [29], SR has generated a lot of interest among vendors, operators, and the research community for a while. On the vendor side, research efforts on SR started around the mid-2000s, with the initial proposal for the SR technology coming from Cisco in 2006, and since then, hardware vendors such as Cisco [54], Juniper [39], and Huawei [35], added SR support in their products. On the operator side, several network operators have publicly announced their use of SR, including Alibaba [61], Bell Canada [70], China Mobile [20], Google [14], Line Corporation [37], Microsoft [53], Vodafone [21], and many others. On the side of the research community, SR also generates a growing interest as shown in Fig. 1, which shows the number of publications that appear, per year, when typing "Segment Routing" on three online scientific libraries: ACM Digital Library, IEEEXplore, and Science Direct. There is a clear rise in the number of publications since 2014, with varied research areas about SR [69]. These studies encompass SR features exploitation to solve classic networking problems, as well as how new functions can be implemented on top of SR to support advanced services.

However, despite all this research, implementation, and standardization efforts, very little is known about the actual SR deployment in the Internet and its actual usage by operators. Padurean et al. [63] have tracked SR with IPv6 as forwarding plane but found no evidence of deployment, probably due to IPv6 Extension Headers filtering [38]. Marechal et al. [52] investigated SR with MPLS as forwarding plane deployment in a large European ISP, relying on data

IMC '25, Madison, WI, USA.

ACM ISBN 979-8-4007-1860-1/2025/10
https://doi.org/10.1145/3730567.3764436

collected by Caida and RIPE Atlas, but the analysis was limited and evidence of Sr deployment was quite low.

Nonetheless, despite publicly available information from operators and our own survey showing that dozens of them actually implement and deploy Sr, the community is, for now, not capable of measuring the presence of Sr. Our paper fills this gap.

In this paper, we present **A**dvanced **Re**velation of **S**egment Routing **T**unnels (AReST), the first methodology for revealing Segment Routing tunnels with Mpls as forwarding plane (Sr-Mpls). To be clear, two forwarding plane implementations are proposed for Sr: Sr over Mpls [12, 27], denoted Sr-Mpls, and Sr over Ipv6 [31], denoted Srv6. Our technique focuses on detecting Sr over Mpls, and detecting Sr over Ipv6 is out of scope. In a nutshell, AReST detects the presence of Sr-Mpls by looking for specific signals of Sr-Mpls implementation found in route traces, which we denote as *flags*. These route traces are collected with `traceroute` augmented with Mpls tunnel data (e.g., Tnt [46, 64]) as well as *fingerprinting* data [3, 50, 68][1] to reveal hardware vendors along a path.

More specifically, our contributions are:

- We survey operators to question their usage of Sr-Mpls. It appears that operators have deployed Sr-Mpls mainly for simplifying Mpls usage (by getting rid of label distribution protocol [5]) and traffic engineering purposes, such as fast reroute.
- We introduce AReST, a new methodology for revealing the presence of Sr-Mpls in `traceroute` paths.
- We apply the AReST methodology towards more than 40 ASes, half of them being known for having deployed Sr-Mpls, while the others were randomly selected from Caida AS rank dataset. In particular, AReST detected Sr-Mpls in 75% of the 20 analyzed ASes that claimed to deploy it. The evaluation was further reinforced by a ground-truth validation involving direct confirmations from network operators – specifically from AS293 (ESnet) – where AReST obtained a perfect precision on their Sr-Mpls deployment.
- We provide a preliminary insight into Sr-Mpls deployment practices. AReST revealed a notable deployment of Sr-Mpls across various ASes, with particularly strong signals within Content, Transit, and Tier-1 providers, as we identified widespread Sr-Mpls adoption in AS8075 (Microsoft) and AS293 (ESnet) that showed the largest proportion of Sr-enabled hops. Our observations further suggest that Sr-Mpls is occasionally used in interworking with a classic Mpls deployment [11], which is coherent with operators' incremental or targeted Sr deployment.

The remainder of this paper is organized as follows: Sec. 2 discusses the required technical background for this paper; Sec. 3 describes our Sr-Mpls survey to operators; Sec. 4 introduces the **A**dvanced **Re**velation of **S**egment Routing **T**unnels (AReST) methodology; Sec. 5 explains how we collected data about Sr-Mpls; Sec. 6 evaluates the AReST methodology, while Sec. 7 provides a first look at Sr-Mpls deployment practices; Sec. 8 positions this paper with respect to the state of the art; finally, Sec. 9 concludes this paper by summarizing its main achievements.

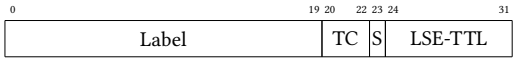

Figure 2: The Mpls label stack entry (Lse) format.

2 Background

In this section, we introduce the required technical background to understand the functioning of Sr-Mpls. In particular, Sec. 2.1 focuses on Mpls, Sec. 2.2 describes how `traceroute` can expose Mpls tunnels, and Sec. 2.3 discusses Segment Routing.

2.1 Mpls

Multi-Protocol Label Switching (Mpls) routers, i.e., *Label Switching Routers* (Lsrs), exchange labeled packets over *Label Switched Paths* (Lsps). In practice, those packets are tagged with one or more *label stack entries* (Lses) inserted between the frame header (data-link layer) and the Ip packet (network layer). Each Lse is made of four fields, as illustrated in Fig. 2: a 20-bit Mpls label used for forwarding the packet to the next Lsr, a 3-bit Traffic Class field for quality of service, priority, and Explicit Congestion Notification [4], a bottom of stack flag bit (to indicate whether the current Lse is the last in the stack [59]), and an 8-bit time-to-live (LSE-TTL) field having the same purpose as the IP-TTL field [1] i.e., avoiding forwarding loops. In an Mpls network, packets are forwarded based on their 20-bit label based on an exact match lookup instead of a longest match lookup on the Ip destination.

Labels may be allocated through the *Label Distribution Protocol* (Ldp[2]) [5]. Each Lsr announces to its neighbors the association between a prefix in its routing table and a 20-bit label it has chosen for a given *Forwarding Equivalent Class* (FEC). Usually, a FEC refers to a set of Ip packets that must be forwarded in the same way by a router. By default, it corresponds to a destination prefix but it may be broader, e.g., packets with same ToS.

A router will announce the same 20-bit label to all its neighbors for a given FEC. It is worth mentioning that the 20-bit label has a *local* significance, i.e., each router independently assigns a 20-bit label for each FEC it handles. Consequently, the label has significance only to the router that announced it, making it unlikely that two routers will assign the same label to the same FEC. This is important information, as if we encounter the same label multiple times, it is a signal that there is more than standard Mpls happening on the path (Sec. 4).

Depending on its location along the Lsp, a Lsr applies one of the three following operations:

- Push. The first Mpls router (*Ingress Label Edge Router* – Ingress Ler) pushes one or multiple Lses in the Ip packet, turning it into an Mpls frame.
- Swap. Within the Lsp, each Lsr makes a label lookup, swaps the incoming label with its corresponding outgoing label, and sends the Mpls packet further along the Lsp.

[1] Fingerprinting refers to the act of dividing network equipment into disjoint classes by analyzing messages sent by that equipment, usually in response to some form of active probing [41, 47].

[2] Labels might also be distributed with RSVP-TE [8] for traffic engineering purposes. Ldp is the most prominent label binding protocol [65, 67] as it is generally the per-default deployment in most Mpls clouds.

- Pop. The *Ending Hop* (EH), the last Lsr of the Lsp, deletes the top Lse. If the Lse was the last in the stack, the Mpls frame is turned back into a standard Ip packet. Otherwise, the remaining stack continues to be forwarded according to the new top of the stack Lse.

As mentioned previously, one can stack one or multiple Lses on an Ip packet for numerous reasons. Inter-domain scalability for transit traffic is the first one [44, 71]. Traffic engineering and fast rerouting are others (with RSVP-TE) [56]. With Ipv6, Lses stacks can be used for 6PE[3] purposes, i.e., either for connecting Ipv6 islands together or using Ldp for Ipv4 to build tunnels carrying both Ipv6 and Ipv4 traffic on dual-stack Mpls routers [66]. Other usages of Lses stacks having multiple labels are *Virtual Private Routing Networks* [34] (VPRN) or *Segment Routing* [29, 32] (Sr).

We have presented here the minimal Mpls background necessary to introduce Segment Routing. Interested readers can find additional details on Mpls in previous work by, e.g., Vanaubel et al. [46, 65].

2.2 `traceroute` and Mpls

Lsrs may send ICMP `time-exceeded` messages when the LSE-TTL expires. If the LSR implements RFC 4950 [15] (as should be the case for all recent OSes), it simply quotes the Mpls Lse stack of the received packet in the ICMP `time-exceeded` message. Such tunnels are called *explicit* [26, 67].

If the first Mpls hop copies the IP-TTL value to the LSE-TTL field rather than setting the LSE-TTL to an arbitrary value such as 255, Lsrs along the Lsp will reveal themselves via ICMP messages, even if they do not implement RFC4950 (in such a case they do not quote the Lse but just reveal their incoming IP address – *opaque* tunnels [26, 67]). Operators can configure this transparency operation using the `ttl-propagate` option provided by the router manufacturer [1]. On the other hand, if the first Mpls router sets the LSE-TTL to 255, the content of the Mpls tunnel will not be revealed through `traceroute` as no TTL will expire along the tunnel (*invisible* tunnels [26, 67]).

Tnt [46, 64], an extension to Paris `traceroute` [7] has been developed to take into account those cases. In addition, Tnt is able to reveal the content of invisible tunnels but without the Lse.

2.3 Segment Routing

Segment Routing (Sr) [12, 29, 30, 32] is a routing technology based on the *source-based routing* paradigm. In this approach, a specific node, called the source, determines the entire path a packet takes through the network by encoding the desired path into a sequence of instructions known as *segments*. This differs fundamentally from classic routing mechanisms where routing decisions occur at every hop, resulting in hop-by-hop path determination. Typically, a segment corresponds to routing instructions such as "forward this packet out of a specific interface" or "forward traffic along the shortest path to a particular router". Multiple segments can exist for the same destination, each specifying different constraints, such as minimizing a metric. Additionally, segments may represent service-based instructions [23] or particular quality-of-service (QOS) treatments.

The Sr data-plane comes in two distinct flavors:

[3] 6PE is an architecture making use of Mpls stack to interconnect Ipv6 islands [25].

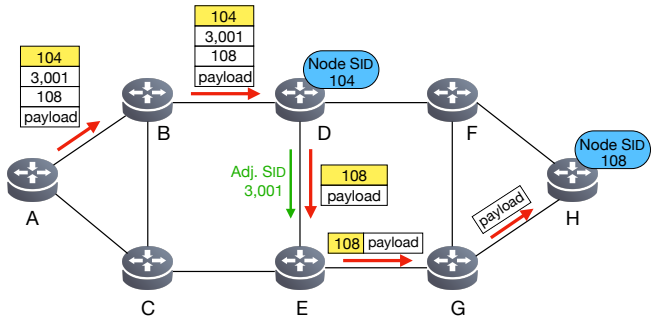

Figure 3: Node SIDs and adjacency SIDs in action. Red arrows indicate the packet forwarding path. Node SIDs 104 and 108 were respectively assigned to routers D and H. Igp link $D \rightarrow E$ is assigned adjacency SID 3,001 by router D. At each step, the top (active) label is highlighted in yellow.

- **Sr-Mpls** [12, 27], which applies Segment Routing directly onto the existing Mpls architecture without altering its forwarding plane. This variant is the primary focus of this paper. In Sr-Mpls, a segment is translated into an Lse. A sequence of segments is encoded as a stack of Lses, the top Lse corresponding to the current segment.
- **Srv6** [31], which uses Ipv6 data-plane through an Ipv6 Extension Header specifically designed for Sr (called the Sr header). Here, each segment is associated with an Ipv6 address, further pushed within the Extension Header.

In Sr, each segment in the network is given a *Segment Identifier* (`SID`). Almost every network entity can be given a `SID`, including: a router (or node), any link with an Igp adjacency on it, unicast prefixes, and anycast prefixes. For network entities supporting dual-stack configurations (e.g., routers), separate `SIDs` are assigned for Ipv4 and Ipv6 respectively.

Segments mainly fall into two categories:

(1) ***prefix* SIDs**: A prefix `SID` is simply a label that instructs routers to forward the packet toward a destination prefix by shortest path.
(2) ***adjacency* SIDs**: An adjacency `SID` is a label that instructs the router to forward packets over a specific Igp link. A router with x `IS-IS` or `OSPF` adjacencies generates exactly x adjacency `SIDs`.

A *node* `SID` is a specific type of prefix `SID` in which the prefix corresponds to the loopback address of a particular node. Using a node `SID` in a segment list guarantees that the packet will be routed along the shortest path to that specific router.

Fig 3 clarifies the above definitions: in this network, router D has a node `SID` of 104 and has also allocated the value 3,001 for its Igp adjacency to router E. Likewise, router H has a node `SID` of 108. Suppose router A wants to send a packet to H but explicitly instructs the path to go through D and then the link $D \rightarrow E$. To do this, A pushes a stack of `SIDs` on the packet: [104; 3,001; 108]. The packet's journey is then as follows: (*i*) B sees the top label 104 and recognizes it as a node `SID` for router D, so it forwards the packet toward D via normal shortest-path routing. When the

packet arrives at D, the top label 104 has fulfilled its function, so D pops that label. (*ii*) Now the label 3,001 becomes the active SID, router D recognizes it as its own Adjacency SID and immediately pops label 3,001 and forwards the packet out its IGP link to E. (*iii*) The next active label is now 108, which is a node SID for router H. Starting from E, all routers on the path to H understand that 108 means "forward toward H" so E and then G will forward the packet along the shortest path to H. When the packet finally reaches H, router H (node SID 108) recognizes the label 108 as referring to itself and pops it, delivering the payload to H.

Furthermore, prefix SIDs have a global significance throughout the domain, while adjacency SIDs have a local significance. It is worth mentioning that "local" and "global" significances do not refer to how the information is advertised but rather to whom the information has a meaning. An SID with global significance has a meaning for all routers within the SR domain: any router can act on that label by forwarding the packet toward the associated prefix. In contrast, an SID with local significance has a meaning only for the node that originated it: only that particular router will perform the forwarding action associated with the label. All routers in the network still learn about every SID, both prefix and adjacency SIDs are distributed to all the routers. In the example above, every router knows that label 3,001 corresponds to D's link to E, but only D will ever forward a packet when 3,001 is the active segment. Meanwhile, label 108 is globally significant, so any router on the way to H can handle it by directing the packet closer to H.

The SR-MPLS control-plane defines how SID information is distributed throughout the network. Unlike classical MPLS, SR-MPLS eliminates the need for label distribution protocols such as LDP by taking advantage of link-state IGP protocols such as IS-IS and OSPF. By extending these protocols to include SR capabilities [9, 10], node and adjacency SIDs are efficiently distributed among network routers. When sharing information, SIDs can be directly provided as absolute MPLS labels or as indexes into an MPLS label space, introducing concepts like *Segment Routing Global Block* (SRGB) (implemented by most vendors) and *Segment Routing Local Block* (SRLB) (implemented by some vendors like Cisco and Huawei).

The SRGB represents a reserved label range from which global node SIDs are allocated. Using this model, node SIDs themselves are not absolute MPLS labels; instead, they serve as indexes for generating 20-bit labels. Similarly, the SRLB serves as a dedicated label range for generating adjacency SIDs. However, implementation varies among vendors; for instance, Juniper does not implement a separate SRLB and instead allocates adjacency SIDs from the dynamic MPLS allocation pool.

Both SRGB and SRLB ranges, when used, are shared across the network through the IGP, allowing nodes within the network to use different ranges from their neighbors. Even though it is permitted, using different SRGB ranges across routers in the same domain is not recommended by RFC 8402 [29] and vendor configuration guidelines. Fig. 4 illustrates such a scenario. A router maps a SID to an MPLS label by adding the SID value to the lowest SRGB value of the subsequent hop toward the destination. Consequently, when routers share identical SRGB ranges, the same label persists across multiple hops (e.g., routers $P5$ and $P6$ both show label 33,007). However, if ranges differ between routers (e.g. $P2$, $P3$, and $P4$), each hop must swap the active label, re-mapping the SID according to the subsequent hop's SRGB range.

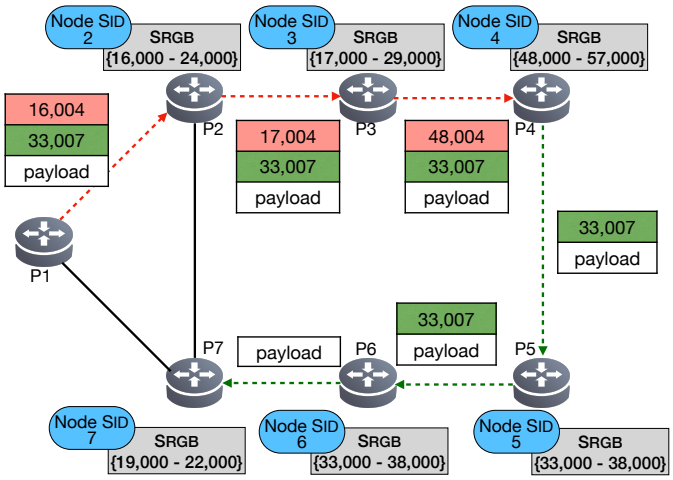

Figure 4: Segment Routing in action. $P1$ pushes a label stack of depth 2 to reach $P7$ by transiting through $P4$: the top label (active label) in red represents the segment to $P4$ and the bottom label in green represents the segment to $P7$.

Table 1: Default vendor-specific Segment Routing Global Block (SRGB) and Segment Routing Local Block (SRLB) MPLS label ranges. Label values in the range 0–255 are reserved for specific MPLS purposes; further details can be found in the relevant RFCs [13, 28, 42, 43, 55, 59].

Label Range	Usage
16,000–23,999	Cisco default SRGB [22]
15,000–15,999	Cisco default SRLB [22]
16,000–47,999	Huawei default SRGB [36]
≥48,000 (user-defined size)	Huawei base SRLB [36]
900,000–965,535	Arista default SRGB [6]
100,000–116,383	Arista default SRLB [6]

Vendor-specific implementations differ in the flexibility offered to users regarding SRGB and SRLB configurations. Some vendors require user-defined ranges, while others use predefined default ranges. Some Segment Routing ranges from well-known vendors are listed in Table 1. In Sec. 4, we will discuss how these configuration variations can help determine the presence of SR-MPLS deployment.

3 Operators Survey

Fig. 1 shows that SR has been an important subject of research during nearly the last ten years, with a peak in 2024. Most of the papers published focus on SR improvements for traffic engineering purposes. However, in practice, little is known about the actual SR (and more specifically, SR-MPLS) deployment and usage by operators.

Table 2: Summary of survey questions and answer types. All questions included an open-ended "Other" option.

Question	Type	Answer Options
What vendor equipment do you use for Sr-Mpls?	Multiple choice	Cisco, Juniper, Huawei, Nokia, Arista, MikroTik, Dell, FreeBSD, Linux, Alcatel, Brocade
If your vendor provides a recommended Srgb, do you follow it?	Closed (yes/no)	Yes, No
If your vendor provides a recommended Srlb, do you follow it?	Closed (yes/no)	Yes, No
Why do you use Sr-Mpls?	Multiple choice	Traffic Engineering, Carry Best Effort Traffic, Simplify Mpls Management, Network Resilience, Carry Traditional Services (e.g., VPNs)

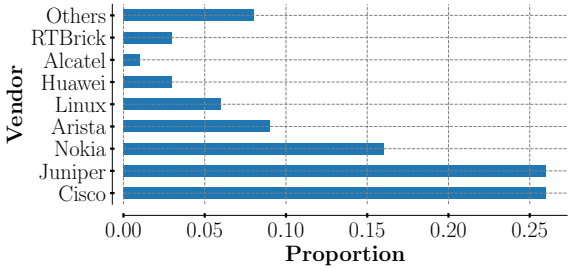

(a) Hardware equipment used for Sr-Mpls.

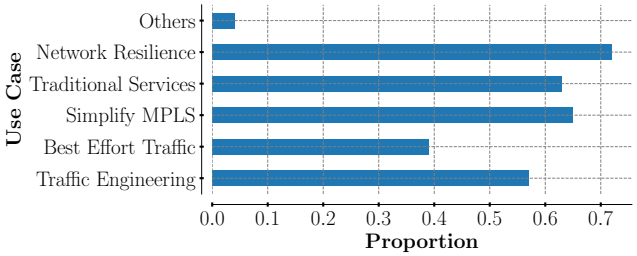

(b) Sr-Mpls usage.

Figure 5: Survey results ($N = 46$). These were multiple-choice questions; thus, percentages do not sum to 100%.

To fill this gap, we conducted a survey among network operators regarding their Sr-Mpls deployment practices. The survey was submitted via the IETF, RIPE, and NANOG mailing lists on January 15th, 2025. At the time of writing this paper, we received 46 responses. The list of questions and associated answers proposed to operators is shown in Table 2

Fig.5 summarizes the results from the survey. First of all, the 46 respondents do deploy Sr-Mpls. Then, as shown in Fig.5a, Cisco and Juniper dominate the vendor market, followed by Nokia, Arista, Linux, and Huawei. Fig. 5b highlights that operators primarily use Sr-Mpls for enhancing network resilience, e.g., fast reroute. Additionally, operators adopt Sr-Mpls to simplify networks by eliminating complexities associated with Mpls and Ldp. As explained in Sec. 2, Sr eliminates the need for Mpls signaling protocols such as Ldp. As a result, operators enable network simplification by removing a protocol from their network. Fig. 5b also indicates that Sr-Mpls is deployed for supporting traditional services, e.g., Virtual Private Networks (VPNs). Interestingly, around 40% of respondents also use Sr-Mpls to transport best-effort traffic, i.e., traffic that is not subject to explicit traffic-engineering mechanisms or strict Quality of Service (QOS) guarantees. We asked whether operators typically modify the Srgb and Srlb ranges when default ranges are provided by vendors (see Table 1). For the Srgb range, 70% keep the recommended range, while 30% prefer customized ranges. For the Srlb, 67% maintain the default range provided and 33% choose to implement custom ranges. Justifications for not considering the default ranges refer to interoperability of different hardware vendors along a path.

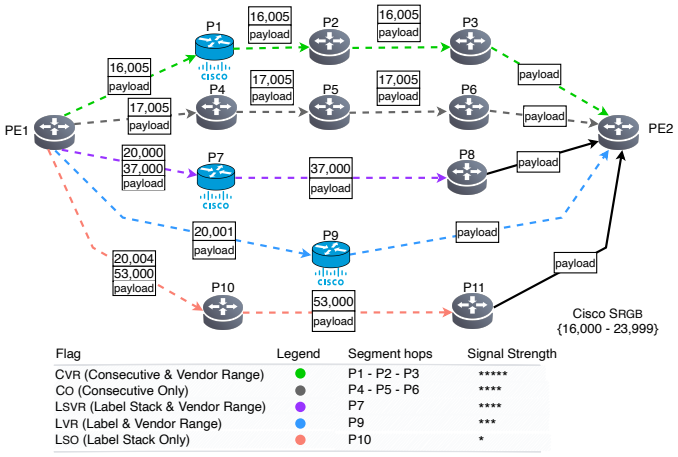

Figure 6: AReST methodology with the Sr-Mpls detection flags. Blue routers are Cisco routers that have been identified through fingerprinting mechanisms [3, 50, 68] while gray routers could not be fingerprinted. For each flag, *PE1* is the Sr source, i.e., the router encoding the Sr path by pushing Lses in the initial Ip packet. Dotted arrows highlight the hops that form the detected Sr-Mpls segment. Vendor's well-known Segment Routing ranges are listed in Table 1.

4 AReST Methodology

This section dives into **A**dvanced **Re**velation of **S**egment Routing **T**unnels (AReST), our detection methodology for Sr-Mpls. AReST operates as a Tnt post-processing tool. It takes as input Paris traceroute [7] paths collected with Tnt [46, 64]. These traces

correspond to standard Paris `traceroute` outputs but, when possible, TNT also exposes the associated MPLS LSE (Sec. 2.2). Then, for each path, and for each hop on a path, AReST uses fingerprinting techniques based on initial TTL signatures [50, 68] and SNMPv3 [3] to reveal the hardware vendors of the hops along a path. The key contribution of AReST lies in parsing these augmented paths and systematically highlighting contiguous portions that exhibit signals of SR-MPLS, which are called segments. Each detected segment is associated with a specific detection flag, and each flag is in turn assigned a signal strength reflecting its reliability and the likelihood of false positives. A segment, in this context, is a contiguous sequence of hops – excluding the source router – that has raised one of our detection flags (see Fig. 6). In the following, we discuss how we designed those flags, ordered by their decreasing signal strengths to reveal the presence of SR-MPLS.

4.1 Consecutive & Vendor Range (CVR)

With classic MPLS (see Sec.2.1), label bindings are local, meaning that each router independently assigns labels for the FECs it manages, taking labels from its own dynamic label pool. However, with SR-MPLS, a segment might be identified by a given 20-bit label, even if that segment is made of multiple hops (see Sec. 2.3). The same 20-bit label will, therefore, appear consecutively in the path as the source node must predefine the entire path, and each hop is advised to use the same starting SR range.

In addition, Table 1 provides a list of well-known 20-bit label values. In particular, it informs about specific label ranges suggested by vendors, e.g., {16,000; 23,999} is the default SRGB range for Cisco devices. We know, thanks to our survey (see Sec. 3, Fig. 5a), that those hardware vendors are present in networks exhibiting SR-MPLS.

Putting together those two key points leads to our first AReST flag: *Consecutive & Vendor Range (CVR)*. The CVR flag is triggered when consecutive identical labels are observed[4] and at least one hop can be mapped to a recognized vendor SR-MPLS range (see Table. 1) through fingerprinting. In this case, the hops sharing the same label are noted as being part of an SR-MPLS segment.

As depicted in Fig. 6 (green path), label 16,005 is first pushed by the source $PE1$ and then consecutively observed across routers $P1$, $P2$, and $P3$, with $P1$ being identified as Cisco and its label is part of Cisco SRGB {16000-23999}, so AReST would label the hops $P1$, $P2$, and $P3$ as part of an SR-MPLS segment.

Regarding false positives, if we assume that each router's dynamic label pool has a size N, then the probability of k consecutive routers independently choosing the same label value would be $\frac{1}{N^{k-1}}$. Taking the example of two Cisco routers, the dynamic label pool typically spans 1,032,575 possible labels [22], making the probability of a coincidence roughly 10^{-6}, which is unlikely to happen. Therefore, the Consecutive & Vendor Range (CVR) flag has the highest signal strength (five stars).

4.2 Consecutive Only (Co)

The CVR flag works as long as we are able to identify the hardware vendor behind a router. It means that the CVR flag is strongly associated with fingerprinting accuracy. However, current fingerprinting techniques' coverage is not perfect and it is not always possible to assign a vendor to a router. Further, as discussed in Sec. 3, hardware vendor label ranges are indicative and nothing prevents operators from changing the default range. Indeed, an operator could be forced to change the default range for interoperability reasons between various vendors in their network.

We therefore propose our second flag, *Consecutive Only (Co)*. It is similar to CVR except that the `traceroute` did not expose at least one hop for which we managed to map the label to a vendor SR-MPLS range. CO works only if we are able to detect consecutive hops using the same 20-bit label value, in which case all the hops with the same label are noted as being part of an SR-MPLS segment.

In Fig. 6 (gray path), routers $P4$, $P5$, and $P6$ consecutively forward packets using label 17,005. Even though 17,005 falls within Cisco SRGB range, these routers could not be fingerprinted, we thus have no explicit confirmation that they are Cisco devices. Consequently, the CO flag is raised, and the hops $P4$, $P5$, and $P6$ are noted as being part of an SR-MPLS segment.

Regarding false positives, the same argument discussed for the CVR flag applies here as well. Therefore, the CVR flag also has a high signal strength (four stars).

4.3 Label Stack & Vendor Range (LSVR)

As discussed in Sec. 2.3, we know that one of the core features of SR is its ability to steer packets along explicitly defined paths by encoding multiple segments in a stack of labels. Encountering an LSE stack thus has the potential to reveal SR-MPLS presence. Indeed, as already mentioned by Vanaubel et al. [66], LSEs stacks are not that frequent in the MPLS landscape: they are mainly used for VPN or traffic engineering purposes. Vanaubel et al. stated that 80% of the tunnels observed (under IPv4) were exhibiting LSEs with a single label depth. Using CAIDA [16] and RIPE Atlas [57] `traceroute` data, we extended their analysis by tracking the evolution of stack sizes from 2015 to 2025 (see Fig. 7). We observed that an LSE stack size greater than 2 is not that frequent: approximately 20% of the traces in the CAIDA dataset (see Fig. 7a) and 10% in the RIPE Atlas dataset (see Fig. 7b) encountered such stack sizes. At this point, it is still not possible to differentiate between a classic MPLS and an SR-MPLS stack, but if the top of the stack LSE encompasses a 20-bit label value belonging to one of the vendor SR ranges (see Table 1), then it is likely that we are in the presence of Segment Routing.

Therefore, our third flag, called *Label Stack & Vendor Range (LSVR)*, is triggered when a hop exhibits an MPLS stack with at least 2 LSEs, and the top label (active label) falls within the recognized vendor SR label range for that router.

Fig. 6 (purple path) illustrates it with router $P7$, identified as Cisco, which exposed a stack [20,000; 37,000] with top label falling within Cisco SRGB. To avoid misclassifications, router $P8$ is not included in the segment, as it may belong to a classic MPLS domain involved in an interworking scenario between SR-MPLS and traditional MPLS. We will discuss such cases in more detail in Sec. 7.

In the context of this flag, a false positive will arise when an operator decides to change the vendor's default SRGB or SRLB and to use that label range for another purpose other than Segment Routing. To assess the likelihood of such false positives, we refer back to our

[4] To handle cases with differing neighbor SRGB ranges, the flag is also triggered if two labels share a common suffix (e.g., 16,005 → 13,005).

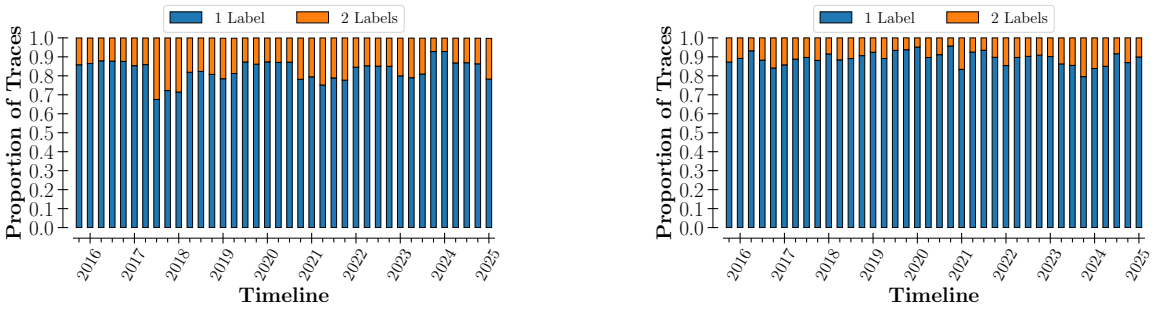

(a) Caida dataset.

(b) RIPE Atlas dataset.

Figure 7: Mpls stack size evolution from December 2015 to March 2025. Data were sampled in March, June, September, and December of each year. `traceroute` data were collected: (*i*) on April 7th, 2025 from the Caida Archipelago measurement platform using three Ark nodes located in the Netherlands, the USA, and Japan [16]; (*ii*) on May 13th, 2025 from the RIPE Atlas platform [57], using Measurements `1835746` (Sweden), `1698335` (USA), and `1404333` (Japan).

operator survey (Sec. 3), indicating that a significant proportion (around 70%) of network operators prefer to retain default vendor-specific Srgb and Srlb label ranges. For the remaining 30%, changes to these default ranges are generally motivated by Sr interoperability concerns within multi-vendor environments. Indeed, RFC 8402 recommends consistent Sr label ranges throughout a domain, but this can be challenging in multi-vendor situations where default vendor ranges inherently differ, encouraging operators to modify their Sr ranges accordingly. In such cases, although the Sr default range may have been changed, it is likely that the interface raising the flag is still involved in an Sr-Mpls multi-vendor environment. For that reason, we keep a high signal strength for that flag (four stars).

4.4 Label & Vendor Range (Lvr)

A specific case of the Lsvr flag occurs when the stack contains only a single Lse, meaning no additional labels are present. We separate this case from Lsvr because, as discussed earlier, stacks of depth greater than 2 are more distinctive of Segment Routing.

Consequently, our fourth flag, called *Label & Vendor Range (Lvr)*, applies when a stack with a single Lse is observed but can still be mapped to the router's known label range.

An example of Lvr flag is provided in Fig. 6 with router $P9$ (blue path).

Concerning false positives, the discussion remains the same as for Lsvr. Due to the absence of an explicit sequence of labels and a stack encompassing several Lses, the Lvr flag signal strength is moderate (three stars).

4.5 Label Stack Only (Lso)

The Lsvr flag works as long as we have an Lses stack with the top of the stack being identified as belonging to a well-known vendor range. As illustrated in Fig. 7, Lses stacks are not that frequent. Therefore, the presence of Lses stack, without the mapping to a well-known vendor range and the consecutive hops, can still be a marker for Sr-Mpls.

Therefore, our last flag, *Label Stack Only (Lso)*, is raised when a single label stack with two or more labels is observed, with no label sequences or Sr vendor range mappings possible.

The Lso flag is illustrated in Fig. 6, at hop $P10$ that presents a stack of depth 2 (orange path).

Regarding false positives, we can exclude the two following cases which would correspond to standard Mpls: Stacks involved in 6PE[3] because we made our measurements under Ipv4, and stacks associated with VPNs as our analysis covers only public Ip addresses. However, for some other cases, we cannot distinguish between classic Mpls and Sr-Mpls, for instance when stacks are used for network resilience purposes, which is one of the main Sr-Mpls usage according to our survey (see Sec. 3). We therefore keep this flag but assign it with the lowest signal strength (one star).

Even though this flag has the lowest signal strength, it remains valuable because advanced Sr mechanisms such as *service* SIDs [23] and *Sr policies* [33] often make use of complex, multi-label stacks. We have indeed captured instances of these advanced Sr features, a deeper discussion of these scenarios is provided later in Sec. 6.

5 Data Collection

Our objective is to detect Segment Routing within an AS, so we need to perform targeted measurements to the ASes of interest to reveal their intra-domain topology, enriched with Mpls tunnel information. Existing datasets such as Caida Archipelago [16] or RIPE Atlas [19] do not suit our needs as they do not explicitly target ASes of interest [16], or do not provide the opportunity to completely reveal tunnels [19]. This section describes how we collected data for investigating Sr-Mpls deployment in 60 ASes, documented in Appendix B.

Among these 60 ASes, there are 25 that, according to private communications with Cisco, have deployed Sr-Mpls. In addition, ten ASes confirmed their Sr-Mpls deployment through our survey (Sec. 3). As such, these 35 will serve as validation cases (Sec. 6.2). Among the surveyed ASes, one went further and manually reviewed all inferences produced by AReST, confirming their accuracy. We therefore consider this AS our ground truth reference (Sec.6.1). The

remaining 25 ASes were selected from the CAIDA AS rank as of October 10th, 2024 [17], considering the ASes with the biggest customer cone [17] in four categories: Stub, Content Providers, Transit, and Tier-1s. For clarity, results in the remainder of the paper will reference each probed AS with its identifier of the form AS#ID, as defined in Table 5. AS identifiers highlighted in red refer to ASes having deployed SR-MPLS according to Cisco. Those in blue represent ASes that confirmed SR-MPLS deployment via our survey (Sec. 3). Those in green denote ASes confirmed by both Cisco and survey responses. Finally, those in black refer to ASes selected from CAIDA AS rank. The ASes identifiers are categorized as follows: #1–12 represent Stub ASes, #13–25 Content Provider ASes, #26–52 Transit ASes, and #53–60 Tier-1 ASes.

Then, for each AS, we built a list of target IP addresses, using ANAXIMANDER [51], an AS mapping framework. In a nutshell, given a set of vantage points (VPs) and a given targeted AS, ANAXIMANDER aims at discovering the most complete map of this AS using the minimum amount of probes to enable a quick and efficient measurement campaign. To do so, ANAXIMANDER starts by collecting RIBs. Afterwards, ANAXIMANDER finds an initial pool of targets expected to transit the AS of interest, applies pruning techniques to this initial pool to reduce the probing load and, finally, sorts and schedules targets in preparation for the probing phase. At the end, ANAXIMANDER produces an ordered list of targets for probing the AS of interest.

In this paper, we ran ANAXIMANDER, on April 1st, 2025, with data collected from 63 geographically distributed BGP collectors from RouteViews [62] and RIPE RIS [58] projects. In the end, ANAXIMANDER generated 60 distinct target lists of IPv4 addresses, one list per AS of interest.

We subsequently probed these ASes using TNT [46, 64], an extension to Paris traceroute [7] designed to reveal MPLS tunnels, whether partially (i.e., only IPv4 addresses of routers involved in an MPLS tunnel) or fully (i.e., the IPv4 address comes with the LSE stack). Measurements with TNT were made from April 2nd to April 6th, 2025, from 50 geographically distributed VPs. Since ANAXIMANDER selects targets based on BGP collectors, vantage points (VPs) used to probe them should ideally be topologically close to those collectors. However, deploying our own measurement infrastructure on the collectors themselves is not feasible, and impractical at scale. We therefore adopt a best-effort approach: we provision VPs from cloud providers that are *geographically close* to the BGP collectors used by ANAXIMANDER. These VPs, including their geographic locations and hosting providers, are detailed in Appendix A.

TNT was configured to send UDP probes for revealing as many links as possible [45]. Those probes were sent at a reasonable pace, i.e., 1,250 pps, with TTL initialized to 1 as we do not know the exact hop count between a VP and the targeted AS. Each VP dealt with the same target list, but shuffled between VPs to avoid appearing as an attack. This approach highlights SR-MPLS tunnels from multiple entry points into the targeted AS.

To accurately identify the entry Autonomous System Border Router (ASBR) of each targeted AS, we relied on bdrmapIT [49] for annotating routers with their respective AS ownership, delimiting the targeted AS from the rest of the Internet. The accuracy of bdrmapIT annotations can be further improved by providing additional router aliasing information, which maps interfaces with specific routers. For this purpose, we relied on MIDAR [40], an IP-ID-based inference technique, complemented by APPLE [48], another router aliasing inference tool.

This methodology resulted in a dataset encompassing 7, 722, 621 traceroutes. As we are interested in SR-MPLS deployed by specific ASes, we restricted the dataset to traces inside each AS of interest, thanks to bdrmapIT. To ensure an efficient coverage, we further filtered out ASes for which ANAXIMANDER failed to reveal a sufficient amount of hops. In particular, we excluded ASes with fewer than 100 distinct IPv4 addresses observed across the 50 VPs. We believe this threshold is the minimum acceptable level to obtain a reasonably representative coverage of an AS: below 100 discovered interfaces, the view is too sparse to make reliable inferences on intra-domain mechanisms such as SR-MPLS. As a result, 19 ASes were excluded: #1, #4–6, #8–12, #18, #21–23, #32, #45, and #48–51. The resulting dataset contains approximately 1.9 million distinct IPv4 addresses. Table 5 reports these statistics in detail, including for each AS the number of discovered IPv4 addresses obtained by our methodology.

To highlight the presence of SR detection flags based on vendor-specific SRGB or SRLB ranges—specifically the CVR, LSVR, and LVR flags introduced in Sec. 4, one needs to map routers to their underlying hardware vendors. This is achieved by using a combination of TTL-based [50, 68] and SNMPv3-based [3] fingerprinting methods. TTL-based fingerprinting is already implemented in TNT. For SNMPv3-based fingerprinting, we considered the latest dataset (dated September 10th, 2024) to annotate our observations.

Regarding flag assignment, the TTL-based fingerprinting method only allows identification of routers as either Cisco or Huawei, but it cannot distinguish between the two vendors because they share the same TTL signature. Therefore, vendor range flags coming from TTL-based detection are raised for MPLS labels within the intersection between Cisco and Huawei SR label space, specifically {16,000; 23,999} (see Table 1). This is as opposed to SNMPv3-based fingerprinting that allows exact differentiation between vendors, thus flags are set based on their respective known SR ranges listed in Table 1. Unfortunately, Arista fingerprinting is not present in the SNMPv3 dataset and thus could not be exploited within our detection methodology. In cases where both methods provide different results for the same hop, SNMPv3-based fingerprinting takes precedence.

6 AReST Evaluation

In this section, we evaluate the effectiveness of the **A**dvanced **Re**velation of **S**egment Routing **T**unnels (AReST) methodology in detecting SR-MPLS deployments. We applied AReST to the 41 ASes of Sec. 5. Our evaluation examines both the prevalence of SR flags in these ASes and the accuracy of AReST inferences. In summary, first, AReST achieved perfect precision on a unique ground truth dataset, provided by one operator who manually confirmed all AReST inferences (Sec. 6.1). In addition, we considered the 20 analyzed ASes claiming to deploy SR-MPLS as validation cases. In this other form of evaluation, AReST was able to detect clear evidence of Segment Routing in 75% of them (Sec. 6.2). Then, for the set of all the ASes, we discuss the different flags found showing evidence of SR-MPLS

Table 3: AReST validation on AS#46 (ESnet) for 17,687 distinct segments flagged as involved in Sr-Mpls.

Flag	Number		Metric (%)			
	Raw	%	TP	TN	FP	FN
Cvr	0	0%	–	–	–	–
Co	16,919	95.6%	100%	–	0%	0%
Lsvr	0	0%	–	–	–	–
Lvr	0	0%	–	–	–	–
Lso	768	4.4%	100%	–	0%	0%

deployment (Sec. 6.2). Finally, we discuss the limits of what can be evaluated with our dataset (Sec. 6.3).

6.1 Ground Truth Validation

We validated the AReST methodology by contacting network operators who responded to our survey. We provided them a list of Ip interfaces, augmented with DNS names, that were identified as involved in Sr-Mpls. From those contacts, only AS#46 (ESnet) replied. A summary of the validation by AS#46 is provided in Table 3. A true positive (TP) corresponds to an Ip interface identified as being Sr-Mpls by AReST and is actually an Sr-Mpls interface, while a false positive (FP) corresponds to an Ip interface identified as being Sr-Mpls by AReST whereas it is only traditional Mpls. The definition of true negatives (TN) and false negatives (FN) can be deduced from the above definitions.

We found 103 distinct Ip interfaces involved in more than 17,000 Sr segments flagged by AReST. 95% of these segments were attributed to the Co flag, while the remaining 5% were Lso-flagged. The operator confirmed the accuracy of the entire provided list, leading to a 0% false positive rate within AS#46 for both Co and Lso flags. In addition, the operator from AS#46 further confirmed that they were running Sr-Mpls everywhere, i.e., without relying on traditional Mpls. In our analysis, no interface within AS#46 was classified as traditional Mpls, which indicates that AReST also presented a false negative rate of 0% within AS#46 in these specific scenarios involving flags Co and Lso.

No Cvr, Lsvr, or Lvr flags were identified within AS#46 because no hops responded to SNMPv3-based or TTL-based fingerprinting.

6.2 Segment Routing in the Wild

At first, it is worth noting that our detection of Sr segments is likely conservative since Mpls tunnels do not consistently reveal Mpls Lses. Indeed, according to Donnet et al. [26], Mpls tunnels can be categorized depending on their implementation of RFC4950 [15][5] and propagation of the IP-TTL field[6] within the tunnel: explicit, implicit, opaque, and invisible types. We have considered all types of tunnels in our analysis, but only explicit tunnels fully expose Mpls Lses, making them eligible for all detection flags described in Sec. 4. Opaque tunnels expose only the last hop Lse, limiting their eligibility to flags Lsvr, Lvr, and Lso. Approximately 76% of the tunnels we observed were explicit (see Fig. 13a in Appendix C for details).

Fig. 8 summarizes the AReST detection results. These detections predominantly occurred among Content, Transit, and Tier-1 ASes. Generally speaking, AReST detected Sr-Mpls in 75% of the 20 analyzed ASes that have claimed to deploy Segment Routing. 60% out of the 75% were mainly identified by the strongest flags (Cvr, Co). The remaining 25% typically lacked sufficient explicit tunnels; for instance, AS#44 (Midco-Net) exhibited explicit tunnels in 5% of its paths, while ASes #2 (Iliad Italy), #3 (NTT Docomo), and #16 (Rakuten) had no explicit tunnels at all (see Appendix C). Combined with our operator validation (see Sec. 6.1), these results underline the good detection coverage provided by AReST.

Moreover, for ASes for which we have no information on whether they have deployed Sr, AReST still detected evidence of Sr-Mpls in 94% of cases. Although initially surprising, nearly one third of these ASes presented at least 90% of Lso-flagged segments. Given that these ASes lack external confirmation of Sr deployment and that Lso carries the lowest signal strength, a more conservative assessment is required when interpreting these findings. In particular, the limitations of such inferences – along with the solutions we adopt to mitigate their impact – will be discussed in Sec. 6.3.

Typically, the Lso flag is the most frequently observed, followed by the strong indicator Co. Significant detections of Co were noted in ASes #13 (Alibaba), #27 (Bouygues), #28 (Bell Canada), and #46 (ESnet), all previously confirmed via Cisco or survey responses. Flags Cvr, Lvr, and Lsvr appear less frequently due to the limited fingerprinting coverage. Indeed, among the subset of identified Sr hops, only 23% were successfully fingerprinted, with 98% of these identifications obtained from TTL-based fingerprinting. As explained in Sec. 5, this results in an inability to differentiate between Cisco and Huawei devices, leading us to restrict our vendor-specific label matching to the overlapping range shared by both vendors, specifically labels 16,000-23,999 (Table1).

ASes #31 (KDDI), #38 (Telecom Italia), #40 (Hurricane Electric), and #55 (Orange) are part of the greatest contributors featuring Cvr, Lsvr, and Lvr flags, mainly due to a higher amount of fingerprinted hops within these ASes. Additionally, for flags Cvr and Co, the AReST methodology implements suffix-based label matching to detect sequences across differing Srgb ranges (Sec. 2.3). However, our observations revealed that only 0.01% of matches were suffix-based, indicating an alignment to RFC 8402 guidelines recommending consistent Srgb ranges among routers.

While flag Lso is difficult to discriminate between classical Mpls and Sr-Mpls, analyzing Lse stack sizes offers further insights. Fig. 9 compares the distribution of stack sizes detected by the strongest Sr flags, i.e., Cvr, Co, Lsvr, and Lvr (Fig. 9a), against those seen in traditional Mpls and Lso contexts (Fig. 9b). The analysis reveals a notably higher tendency for stack sizes ≥ 2 in Sr contexts, with such stacks appearing approximately 20% more frequently on average. The latter provides further confirmation of our motivation for designing the flag Lsvr (see Sec. 4). However, ASes #46 (ESnet), and #52 (Execulink) exhibited an important amount of stacks ≥ 2 regardless of the context. A deeper look indicates that ASes #46 and #52 employ unshrinking stacks, less frequent in traditional Mpls but plausible in advanced Sr scenarios involving service SIDs [23]

[5] If the router implements RFC4950, as it should be the case for all recent OSes, it simply quotes the Mpls Lse stack of the received packet in the ICMP time-exceeded message.

[6] The Ingress router, that pushes the very first Lse on the packet, may decide to copy the IP-TTL into the LSE-TTL – see Fig. 2 – rather than setting the LSE-TTL to an arbitrary value such as 255. If so, routers inside the Mpls tunnels will reveal themselves via ICMP messages, even if they do not implement RFC4950.

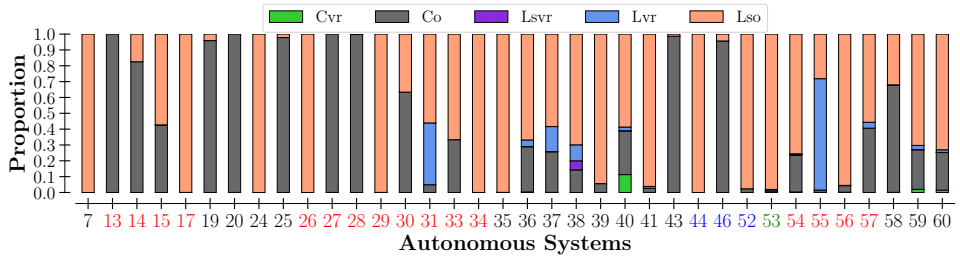

Figure 8: Proportion of SR segments flagged by each AReST detection flag. AS identifiers match those listed in Table 5. Colors indicate source of SR-MPLS confirmation: red (Cisco), blue (survey), green (both), and black (no explicit confirmation). Identifier ranges reflect AS roles: #1–12 Stub, #13–25 Content Providers, #26–52 Transit, #53–60 Tier-1.

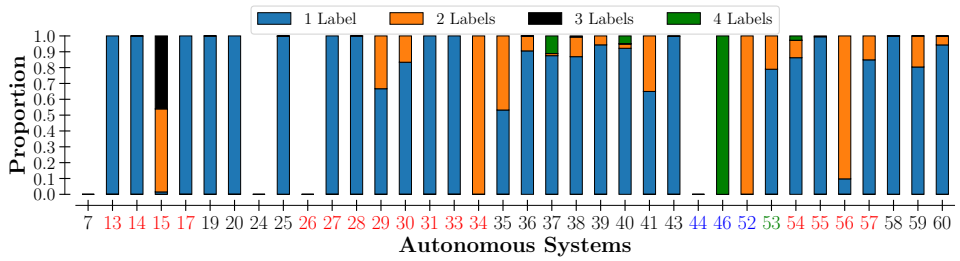

(a) Distribution of observed LSE stack sizes in segments flagged with strong SR-MPLS flags: CVR, CO, LSVR, and LVR.

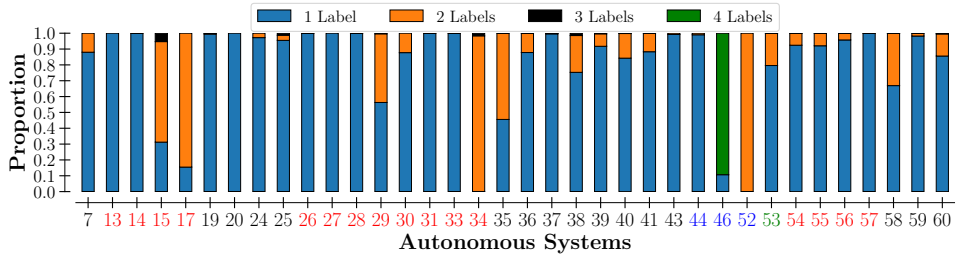

(b) Distribution of observed LSE stack sizes across traditional MPLS hops and LSO-flagged hops.

Figure 9: Distribution of LSE stack sizes observed in the dataset. AS identifiers match those listed in Table 5. Colors indicate source of SR-MPLS confirmation: red (Cisco), blue (survey), green (both), and black (no explicit confirmation). Identifier ranges reflect AS roles: #1–12 Stub, #13–25 Content Providers, #26–52 Transit, #53–60 Tier-1.

or SR policies [33]. Service SIDs use multi-layer stacks to separate transit and service delivery segments, often resulting in unshrinking stacks observable at the destination. Furthermore, SR policies allow one hop on a path to dynamically replace certain SIDs with new, potentially deeper, stacks. It should also be noted that the LSO segments of AS#46 (ESnet) have been explicitly confirmed by the ground truth (see Sec. 6.1), so the characteristics observed on the stack sizes are unlikely to be a coincidence. Although intriguing, a deeper exploration of these advanced SR mechanisms falls beyond this paper's scope and represents potential future work. Nonetheless, these observations suggest that the LSO flag is capable of capturing specific SR deployment scenarios.

6.3 Limits

Our detection methodology relies on a set of flags characterized by varying degrees of signal strength, as defined in Sec. 4. Strong flags such as CVR and CO are based on explicit label sequences, or SR label ranges confirmation. However, flags like LSO inherently have lower signal strength, increasing the risk of misclassification.

For ASes whose SR-MPLS deployment status is externally confirmed (e.g., via survey responses), the strength of flag-based detections is further reinforced. However, when assessing ASes without external confirmations, our evaluation relies entirely on the signals provided by the AReST flags.

Even though the ground truth validation (Sec. 6.1) and our analysis (Sec. 6.2) suggest that the LSO flag can correctly identify specific

Sr-Mpls segments, its broader nature, discussed in Sec. 4, inherently introduces the risk of false positives. However, in cases such as ASes #14 (Google), #19 (Amazon), or our ground truth case AS #46 (ESnet), where the Lso flag occurs alongside strong indicators (Cvr, Co, Lsvr, Lvr), the detection strength of Lso-flagged segments is significantly enhanced because explicit evidence of Sr-Mpls has already been confirmed within these networks. On the other hand, scenarios like AS #7 (Proximus), where 100% of the flagged segments are exclusively Lso, remain ambiguous and thus need more cautious interpretation. To maintain a conservative assessment of Sr, segments flagged by Lso will therefore be excluded from further analysis in the remainder of this paper.

Additionally, while false negatives have been shown to be minimal for our Co flag (Sec. 6.1), the potential for false negatives remains in other scenarios. In particular, if a Sr segment consists only of a single hop whose label does not fall within known vendor-specific Sr ranges, our methodology might fail to raise any flags, as such a case cannot be distinguished from classic Mpls behavior. Consequently, specific single-hop Sr-Mpls segments might remain undetected.

7 Preliminary Sr-Mpls Characterization

This section provides a preliminary overview of how Sr-Mpls is deployed and used in practice across the studied ASes. In summary, we found that Segment Routing is primarily present in large Transit, Content, and Tier-1 ASes, while it is almost absent in small Stub networks. We also observed that most Sr tunnels are fully Sr-based, but a notable minority (around 10%) are hybrid Sr tunnels that interwork with Ldp. In these interworking scenarios, an Sr domain typically connects or spans across one or more smaller Ldp domains. In the following, we first quantify the extent of Sr deployment within each AS (Sec. 7.1), then analyze how Sr and Ldp are combined when they coexist (Sec. 7.2).

7.1 Segment Routing Deployment

This section provides a first overview of Sr-Mpls usage by operators. In particular, Fig. 10a illustrates the proportion of traces encountering Sr-Mpls, classic Mpls, and Ip areas, indicating how frequently one might encounter Sr-Mpls along a given path inside an AS. Results show that the contributors to Sr-Mpls in our dataset are mainly found among Content, Transit, and Tier-1 ASes. This is particularly true for ASes #15 (Microsoft), #28 (Bell Canada), #46 (ESNet), and #58 (Arelion), where more than 50% of the targeted traces have hit an Sr-Mpls area. In contrast, Stub ASes in our dataset display almost no Segment Routing deployment, this is mainly due to the fact that only 26% of tunnels observed within Stub ASes were explicit tunnels (see Appendix C).

However, Fig. 10a alone does not fully quantify actual Sr-Mpls usage since even a trace traversing a very short Sr area – such as a single hop – followed by extensive Mpls or Ip regions would still be classified as Sr. To complement this perspective, we present Fig. 10b, which quantifies the number of distinct Ip interfaces involved in each of these three routing mechanisms. This analysis shows that the number of Sr-enabled interfaces is significantly lower compared to the total number of observed Ips within the ASes. Indeed, for 88% of the analyzed ASes, the proportion of Sr-related interfaces represents 10% or less. Notable exceptions to this trend include AS #15 (Microsoft) and AS #46 (ESnet), where approximately 50% and 33% of the observed Ip interfaces were involved in Sr-Mpls, respectively.

We note that these figures represent conservative estimates of Sr deployment due to our strict criteria, which only rely on flags Cvr, Co, Lsvr, and Lvr to identify Sr segments. As discussed in Sec. 6, this ensures a low false positive rate while maintaining a high true positive rate. Additionally, quantifying Sr deployment depends not only on Anaximander's ability to explore the full topology of an AS, but also on the presence of non-explicit tunnels, which inherently limit our measurement coverage. On average, our methodology allowed us to explore approximately 46,000 distinct Ip addresses per analyzed AS.

7.2 Sr-Ldp Interworking

The AReST methodology not only allows us to estimate Sr deployment along a path, but also to characterize its usage practices. In particular, we identified two types of Sr tunnels: (*i*) tunnels that fully expose Sr-Mpls hops are called *full-Sr tunnels*; (*ii*) hybrid tunnels that expose not only Sr-Mpls hops but also classic Mpls hops. The latter category runs Sr along part of the path but also relies on the Label Distribution Protocol (Ldp) for other parts (see Sec. 2). Such tunnels are known as *interworking tunnels* [11]. In our dataset, 10% of the tunnels we observed were involved in interworking scenarios, while 90% were full-Sr tunnels.

Given that a significant minority (10%) of the observed Sr tunnels are involved in an interworking situation with Ldp, we analyzed the different patterns of how Sr and Ldp segments are chained together in these hybrid tunnels. Fig. 11 shows the proportion of the various interworking combinations we observed. Sr to Ldp interworking is the most common scenario we observed, 95% of the tunnels were using that mode. This case arises when the destination of the tunnel is not Sr-enabled, and the Sr hops thus have no prefix SID for that destination. In such cases, an additional control plane component is required: the *Mapping Server* [11]. Its role is to advertise prefix SIDs on behalf of non-Sr-capable routers, such that every Sr router knows how to reach any Ldp router. Ldp to Sr case is more uncommon according to our results (approximately 2%). In the latter case, interworking happens seamlessly, i.e., without any additional signaling or per-flow state required. Every Sr router adjacent to the Ldp domain is responsible for generating Ldp label bindings that mirror the node SIDs they have learned from the Sr domain. These bindings are then advertised via Ldp to their downstream neighbors, ensuring connectivity between the two domains. More advanced chaining patterns, such as Ldp-Sr-Ldp and Sr-Ldp-Sr, were also observed in 2% and 1% of the interworking cases, respectively. These scenarios can be seen as combinations of the simpler Sr-Ldp and Ldp-Sr transitions discussed earlier. As such, their operation follows the same principles at each boundary between domains.

To further qualify the nature of these interworking tunnels, Fig. 12 compares the sizes of the Ldp and Sr clouds. We see that Ldp clouds tend to be smaller, whereas Sr clouds are typically larger. This suggests a deployment scenario where smaller Ldp islands are being interconnected by larger Segment Routing clouds.

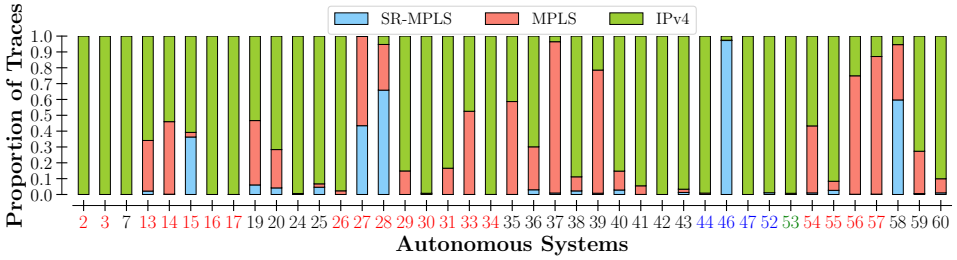

(a) Proportion of traces hitting SR-MPLS, MPLS, and IP areas.

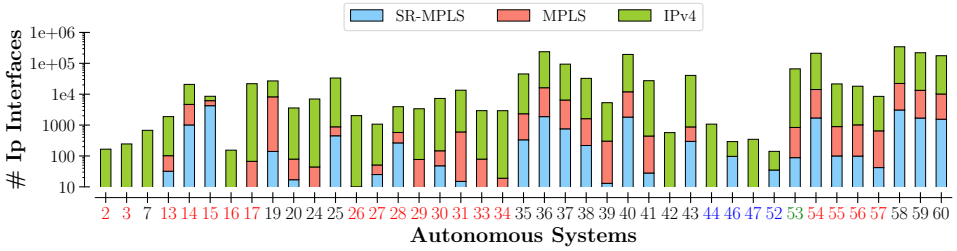

(b) Amount of distinct IP interfaces identified in each areas.

Figure 10: Global view of Segment Routing deployment relative to traditional IP and MPLS areas. Strong SR flags (CVR, CO, LSVR, LVR) are used to identify SR-MPLS areas. AS identifiers match those listed in Table 5. Colors indicate source of SR-MPLS confirmation: red (Cisco), blue (survey), green (both), and black (no explicit confirmation). Identifier ranges reflect AS roles: #1–12 Stub, #13–25 Content Providers, #26–52 Transit, #53–60 Tier-1.

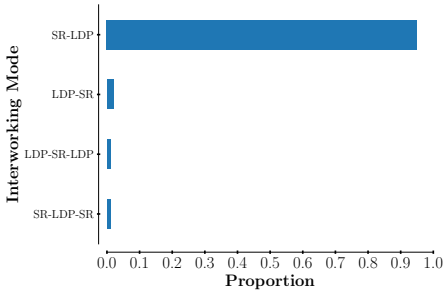

Figure 11: Proportion of the different Interworking modes observed.

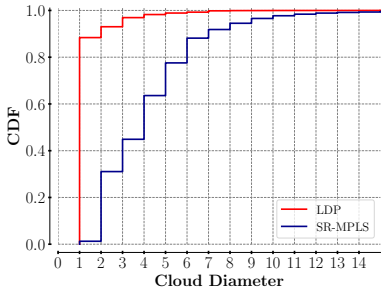

Figure 12: Distribution of LDP and SR cloud sizes involved in Interworking tunnels.

Finally, it is worth noting that despite the demonstrated accuracy of AReST (Sec.6.1), the methodology remains susceptible to a few potential false-negative scenarios. These edge cases, which might lead to misclassifying full-SR tunnels as interworking, have been carefully examined in Sec.6.3.

8 Related Work

Segment Routing has witnessed growing attention in recent years, driven by significant standardization efforts and an increasing number of research activities exploring its diverse use cases. This trend is illustrated in Fig. 1, which shows a clear rise in SR-related publications. Ventre et al. [69] surveyed existing research activities around SR, highlighting applications in monitoring, traffic engineering, failure recovery, path encoding, and performance evaluation, among others.

However, research addressing SR deployment remains limited. One of the first comprehensive deployment studies was conducted by Padurean et al. [63], focusing on the SRV6 data plane. Their work involved analyzing publicly available route collector data for evidence of SRV6 deployments, but failed to detect any clear indications of its use. They then conducted a large-scale `traceroute` measurement campaign, but no SRV6 was observed, not even for ISPs that claimed to have deployed it. Upstream filtering of IPV6 extension headers [38] might explain this absence of observed activity.

For the Sr-Mpls data plane, Marechal et al. [52] processed a `traceroute` dataset and searched for Sr-Mpls signs within Vodafone ISP. Their methodology involved identifying Cisco routers through TTL-based fingerprinting and subsequently mapping observed labels to Cisco's known Srgb ranges. However, their analysis is incomplete compared to this paper, in particular by not taking 20-bit label sequences into consideration.

More generally speaking, Mpls tunnels discovery has been the subject of several research studies during the last fifteen years. In particular, Sommers et al. [60] examined the characteristics of Mpls deployments that are explicitly identified with `traceroute`, as observed in Caida's topology data. Donnet et al. [26] proposed the classification of Mpls tunnels according to the relationship between Mpls and `traceroute`. This classification has been later extended by Vanaubel et al. [64], who also proposed techniques for inferring and revealing Mpls tunnels hidden from classic `traceroute` [67].

In parallel with discovery mechanisms, several studies have looked at the Mpls deployment. For instance, Vanaubel et al. have investigated the Mpls usage in Ipv6 [66]. Al-Qudah et al. [2] have looked at the interaction between Mpls tunnels and path stability. Davila Revelo et al. [24] questioned the relationships between Mpls tunnels and ASes topology. Finally, Vanaubel et al. [65] considered 20-bit labels for identifying transit path diversity inside ASes. However, Vanaubel et al. did not provide any insight 20-bit label values and did not discuss advanced usage of Mpls, such as Segment Routing.

9 Conclusion

This paper discussed Segment Routing. While Sr has been a subject of numerous studies during the last ten years, little is known about its actual deployment by operators. This is exactly the focus of this paper, in particular by considering Sr with Mpls as data plane.

We introduced **A**dvanced **Re**velation of **S**egment Routing **T**unnels (AReST), a methodology for detecting Sr-Mpls. AReST is lightweight as it relies only on `traceroute`-like data augmented with fingerprinting for identifying hardware vendors along a path. We applied AReST towards more than 40 ASes based on data collected from 50 geographically dispersed VPs.

We have shown that AReST is able to reveal a notable deployment of Sr-Mpls across various types of ASes, with particularly strong signals within Content, Transit, and Tier-1 providers. In particular, AReST detected Sr-Mpls in 75% of the 20 analyzed ASes that claimed to deploy it. We also provided preliminary insights into Sr-Mpls deployment practices. Our observations suggest that Sr-Mpls is occasionally used in interworking with a classic Mpls deployment.

Furthermore, direct validation from network operators of AS293 (ESnet) reinforced the robustness of our methodology, confirming the presence of Sr-Mpls where AReST predicted it.

Future work plans to focus on more detailed characterization studies of advanced Sr-Mpls features, as well as longitudinal analyses to track the evolution of Sr-Mpls adoption patterns over time.

Reproducibility

Software developed and used in this paper, as well as data collected, are shared with the community.

Our code for AReST, as well as code developed to test AReST on a controlled environment, is available at https://gitlab.uliege.be/segment-routing.

The data collected and analyzed throughout this paper is available at https://doi.org/10.58119/ULG/AN7GB7.

Acknowledgments

This work has been supported by the CyberExcellence project funded by the Walloon Region, under number 2110186, as well as Les Semestres Thématiques funded by the Walloon Region.

References

[1] P. Agarwal and B. Akyol. 2003. *Time-to-Live (TTL) Processing in Multiprotocol Label Switching (MPLS) Networks.* RFC 3443. Internet Engineering Task Force.
[2] Z. Al-Qudah, M. Alsarayreh, I. Jomhawy, and M. Rabinovich. 2016. Internet Path Stability: Exploring the Impact of MPLS Deployment. In *Proc. IEEE Global Communications Conference (GLOBECOM).*
[3] T. Albakour, O. Gasser, R. Beverly, and G. Smaragdakis. 2021. Third Time's Not a Charm: Exploiting SNMPv3 for Router Fingerprinting. In *Proc. ACM Internet Measurement Conference (IMC).*
[4] L. Andersson and R. Asati. 2009. *Multiprotocol Label Switching (MPLS) Label Stack Entry: EXP Field Renamed to Traffic Class Field.* RFC 5462. Internet Engineering Task Force.
[5] L. Andersson, I. Minei, and T. Thomas. 2007. *LDP Specification.* RFC 5036. Internet Engineering Task Force.
[6] Arista Networks. 2025. EOS 4.33.2F User Manual. https://www.arista.com/en/um-eos/ [Last Access: April 4th, 2025].
[7] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. 2006. Avoiding Traceroute Anomalies with Paris Traceroute. In *Proc. ACM Internet Measurement Conference (IMC).*
[8] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, and G. Swallow. 2001. *RSVP-TE: Extensions to RSVP for LSP Tunnels.* RFC 3209. Internet Engineering Task Force.
[9] A. Bashandy, B. Decraene, C. Filsfils, L. Ginsberg, G. Hannes, and S. Previdi. 2019. *OSPF Extensions for Segment Routing.* RFC 8665. Internet Engineering Task Force.
[10] A. Bashandy, C. Filsfils, S. Previdi, B. Decraene, and S. Litkowski. 2019. *IS-IS Extensions for Segment Routing.* RFC 8667. Internet Engineering Task Force.
[11] A. Bashandy, C. Filsfils, S. Previdi, B. Decraene, and S. Litkowski. 2019. *Segment Routing MPLS Interworking with LDP.* RFC 8661. Internet Engineering Task Force.
[12] A. Bashandy, C. Filsfils, S. Previdi, B. Decraene, S. Litkowski, and R. Shakir. 2019. *Segment Routing with the MPLS Data Plane.* RFC 8660. Internet Engineering Task Force.
[13] M. Bocci, M. Vigoureux, and S. Bryant. 2009. *MPLS Generic Associated Channel.* RFC 5586. Internet Engineering Task Force.
[14] A. Bogdanov. 2019. Introducing Centralized TE.
[15] R. Bonica, D. Gan, D. Tappan, and C. Pignataro. 2007. *ICMP Extensions for Multiprotocol Label Switching.* RFC 4950. Internet Engineering Task Force.
[16] CAIDA. 2024. The CAIDA UCSD IPv4 Routed /24 Topology Dataset, 2014–2024. https://www.caida.org/catalog/datasets/ipv4_routed_24_topology_dataset/ [Last Access: April 7th, 2025].
[17] CAIDA. 2025. *AS-Rank.* https://asrank.caida.org [Last Access: March 1st, 2025].
[18] CAIDA. 2025. *AS Relationships.* https://www.caida.org/catalog/datasets/as-relationships/ [Last Access: March 1st, 2025].
[19] RIPE Network Coordination Center. 2010. Atlas. See https://atlas.ripe.net.
[20] China Mobile Research Institute. 2020. China Mobile Technical White Paper on G-SRv6. https://www.ipv6plus.net/resources/G-SRv6/China_Mobile_Technical_White_Paper_on_G-SRv6.pdf [Last Access: May 9th, 2025].
[21] Cisco. 2018. Cisco and Vodafone showcase Mobile Transport Networking Advancements Via Segment Routing at Mobile World Congress. https://investor.cisco.com/news/news-details/2018/Cisco-and-Vodafone-showcase-Mobile-Transport-Networking-Advancements-Via-Segment-Routing-at-Mobile-World-Congress/default.aspx [Last Access: May 9th, 2025].
[22] Cisco. 2021. *Segment Routing Configuration Guide for Cisco ASR 9000 Series Routers.* Cisco Press, Chapter Chap. 5.
[23] F. Clad, X. Xu, Filsfils C., D. Bernier, C. Li, B. Decraene, S. Ma, C. Yadlapalli, W. Henderickx, and S. Salsano. 2025. *Service Programming with Segment Routing.* Internet Draft (Work in Progress) draft-ietf-spring-sr-service-programming-11. Internet Engineering Task Force.
[24] G. Davila Revelo, M. A. Ricci, B. Donnet, and J. I. Alvarez-Hamelin. 2016. Unveiling the MPLS Structure on Internet Topology. In *Proc. IFIP Network Traffic Measurement and Analysis Conference (TMA).*
[25] J. De Clercq, D. Ooms, F. Le Faucheur, and S. Prevost. 2007. *Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE).* RFC 4798. Internet Engineering Task Force.

[26] B. Donnet, M. Luckie, P. Mérindol, and J.-J. Pansiot. 2012. Revealing MPLS Tunnels Obscured from Traceroute. *ACM SIGCOMM Computer Communication Review* 42, 2 (April 2012), 87–93.

[27] A. Farrel et al. 2017. Segment Routing: Cutting Through the Hype and Finding the IETF's Innovative Nugget of Gold. *IETF Journal* 13, 1 (July 2017).

[28] A. Farrel, S. Bryant, and J. Drake. 2019. *An MPLS-Based Forwarding Plane for Service Function Chaining.* RFC 8595. Internet Engineering Task Force.

[29] C. Filfils, S. Previdi, L. Ginsberg, B. Decraene, S. Litkowski, and R. Shakir. 2018. *Segment Routing Architecture.* RFC 8402. Internet Engineering Task Force.

[30] C. Filsfils, P. Camarillo, J. Leddy, D. Voyer, S. Mathsushima, and Z. Li. 2021. *Segment Routing over IPv6 (SRv6) Network Programming.* RFC 8986. Internet Engineering Task Force.

[31] C. Filsfils, D. Dukes, S. Previdi, J. Leddy, S. Matsushima, and D. Voyer. 2020. *IPv6 Segment Routing Header (SRH).* RFC 8754. Internet Engineering Task Force.

[32] C. Filsfils, N. Kumar Nainar, J. C. Cardona, and P. Francois. 2015. The Segment Routing Architecture. In *Proc. IEEE Global Communications Conference (GLOBECOM).*

[33] C. Filsfils, K. Talaulikar, J. Leddy, D. Voyer, A. Bogdanov, and P. Mattes. 2022. *Segment Routing Policy Architecture.* RFC 8986. Internet Engineering Task Force.

[34] B. Gleeson, A. Lin, J. Heinanen, G. Armitage, and A. Malis. 2000. *A Framework for IP Based Virtual Private Networks.* RFC 2764. Internet Engineering Task Force.

[35] Huawei. 2019. Segment Routing. https://support.huawei.com/enterprise/en/doc/EDOC1100092117 [Last Access: May 9th, 2025].

[36] Huawei. 2025. NE40E Label Space in Segment Routing MPLS. https://support.huawei.com/enterprise/fr/doc/EDOC1100092116 [Last Access: April 4th, 2025].

[37] H. Ichiara. 2019. LINE Data Center Networking with SRv6. https://www.segment-routing.net/images/20190920-LINE_Data_Center_Networking_with_SRv6.pdf [Last Access: May 9th, 2025].

[38] J. Iurman and B. Donnet. 2025. The Razor's Edge: IPv6 Extension Headers Survivability. In *Proc. Passive and Active Measurement Conference (PAM).*

[39] Juniper Networks. [n. d.]. What is Segment Routing? https://www.juniper.net/us/en/research-topics/what-is-segment-routing.html [Last Access: May 9th, 2025].

[40] K. Keys, Y. Hyun, M. Luckie, and kc Claffy. 2013. Internet-scale IPv4 alias resolution with MIDAR. *IEEE/ACM Transactions on Networking (ToN)* 21, 2 (April 2013), 383–399.

[41] T. Kohno, A. Broido, and kc claffy. 2005. Remote Physical Device Fingerprinting. *IEEE Transactions on Dependable and Secure Computing* 2, 2 (May 2005), 93–108.

[42] K. Kompella, L. Andersson, and A. Farrel. 2014. *Allocating and Retiring Special-Purpose MPLS Labels.* RFC 7274. Internet Engineering Task Force.

[43] K. Kompella, J. Drake, S. Amante, W. Henderickx, and L. Yong. 2012. *The Use of Entropy Labels in MPLS Forwarding.* RFC 6790. Internet Engineering Task Force.

[44] J.-L. Le Roux, J.-P. Vasseur, and J. Boyle. 2005. *Requirements for Inter-Area MPLS Traffic Engineering.* RFC 4105. Internet Engineering Task Force.

[45] M. Luckie, Y. Hyun, and B. Huffaker. 2008. Traceroute Probe Method and Forward IP Path Inference. In *Proc. ACM Internet Measurement Conference (IMC).*

[46] J.-R. Luttringer, Y. Vanaubel, P. Mérindol, J.-J. Pansiot, and B. Donnet. 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.

[47] G. F. Lyon. 2009. *Nmap Network Scanning: The Official Nmap Project Guide to Network Discovery and Security Scanning.* Nmap Project. See http://nmap.org/book/toc.html.

[48] A. Marder. 2020. *APPLE: Alias Pruning by Path Length Estimation.*

[49] A. Marder, M. Luckie, A. Dhamdhere, B. Huffaker, K. Claffy, and Jonathan M. Smith. 2018. Pushing the Boundaries with bdrmapIT: Mapping Router Ownership at Internet Scale. In *Proc. ACM Internet Measurement Conference (IMC).*

[50] E. Marechal and B. Donnet. 2020. Network Fingerprinting: Routers under Attack. In *Proc. International Workshop on Traffic Measurements for Cybersecurity (WMTC).*

[51] E. Marechal, P. Mérindol, and B. Donnet. 2022. ISP Probing Reduction with Anaximander. In *Proc. Passive and Active Measurement Conference (PAM).*

[52] E. Marechal, Y. Shao, M. Bruyère, and B. Donnet. 2022. A First-Look at Segment Routing Deployment in a Large European ISP. In *Proc. ACM Internet Measurement Conference (IMC) – poster session.*

[53] P. Mattes. 2016. Traffic Engineering in a Large Network with Segment Routing. https://www.youtube.com/watch?v=CDtoPGCZu3Y [Last Access: May 9th, 2025].

[54] P. S. Mule, J. Mira, and O. Munoz Cueva. 2023. Segment Routing Implementation in Action. https://www.ciscolive.com/c/dam/r/ciscolive/global-event/docs/2023/pdf/BRKMPL-2147.pdf [Last Access: May 9th, 2025].

[55] H. Ohta. 2002. *Assignment of the OAM Alert Label for Multiprotocol Label Switching Architecture (MPLS) Operation and Maintenance (OAM) Functions.* RFC 3429.

[56] P. Pan, G. Swallow, and A. Atlas. 2005. *Fast Reroute Extensions to RSVP-TE for LSP Tunnels.* RFC 4090. Internet Engineering Task Force.

[57] RIPE NCC. 2025. https://atlas.ripe.net [Last Access: May 13th, 2025].

[58] RIPE NCC. 2025. Routing Information Service (RIS). https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris [Last Access: March 31th, 2025].

[59] E. Rosen, D. Tappan, G. Fedorkow, Y. Rekhter, D. Farinacci, T. Li, and A. Conta. 2001. *MPLS Label Stack Encoding.* RFC 3032. Internet Engineering Task Force.

[60] J. Sommers, B. Eriksson, and P. Barford. 2011. On the Prevalence and Characteristics of MPLS Deployments in the Open Internet. In *Proc. ACM Internet Measurement Conference (IMC).*

[61] Y. Su. 2022. Alibaba: Full stack SRv6 towards a Predictable Network.

[62] University of Oregon. 2025. RouteViews Project. http://www.routeviews.org/ [Last Access: March 31th, 2025].

[63] O. V.-A Padurean, Gasser, R. Bush, and A. Feldmann. 2022. SRv6: Is There Anybody Out There?. In *Proc. International Workshop on Traffic Measurements for Cybersecurity (WTMC).*

[64] Y. Vanaubel, J.-R. Luttringer, P. Mérindol, J.-J. Pansiot, and B. Donnet. 2019. TNT, Watch me Explode: A Light in the Dark for Revealing MPLS Tunnels. In *Proc. IFIP Network Traffic Measurement and Analysis Conference (TMA).*

[65] Y. Vanaubel, P. Mérindol, J.-J. Pansiot, and B. Donnet. 2015. MPLS Under the Microscope: Revealing Actual Transit Path Diversity. In *Proc. ACM Internet Measurement Conference (IMC).*

[66] Y. Vanaubel, P. Mérindol, J.-J. Pansiot, and B. Donnet. 2016. A Brief History of MPLS Usage in IPv6. In *Proc. Passive and Active Measurement Conference (PAM).*

[67] Y. Vanaubel, P. Mérindol, J.-J. Pansiot, and B. Donnet. 2017. Through the Wormhole: Tracking Invisible MPLS Tunnels. In *Proc. ACM Internet Measurement Conference (IMC).*

[68] Y. Vanaubel, J.-J. Pansiot, P. Mérindol, and B. Donnet. 2013. Network Fingerprinting: TTL-Based Router Signature. In *Proc. ACM Internet Measurement Conference (IMC).*

[69] P.-L. Ventre, S. Salsano, M. Polverini, A. Cianfrani, A. Abdelsalam, C. Filsfils, P. Camarillo, and F. Clad. 2020. Segment Routing: A Comprehensive Survey of Research Activities, Standardization Efforts, and Implementation Results. *IEEE Communications Surveys & Tutorials* 23, 1 (November 2020), 182–221.

[70] D. Voyer. 2022. Motivation to Migrate from SR-MPLS to SRv6 uSID. Review of the Security Model.

[71] R. Zhang and J.-P. Vasseur. 2005. *MPLS Inter-Autonomous System (AS) Traffic Engineering (TE) Requirements.* RFC 4216. Internet Engineering Task Force.

A Vantage Points

Table 4 shows the list of vantage points (VPs) considered in our measurement campaign. Our VPs are spread over 28 countries scattered all over the world: 30% of the VPs are located in Europe, 34% in North America, 2% in South America, 4% in Africa, 26% in Asia and, finally, 4% in Oceania. Those VPs are owned by four different cloud providers: Amazon AWS (13 VPs), Digital Ocean (1 VP), Google Cloud (21 VPs), and Vultr (15 VPs)

Each VP corresponds to a virtual machine (VM) on which we deployed Tnt. All VMs in our dataset have the same configuration: 2 CPU cores and 4GB of RAM.

Table 4: Virtual machines used as vantage points for our measurement campaign.

ID	Cloud Provider		City/State	Country
	Name	ASN		
VM1			Tokyo	Japan
VM2			Seoul	South Korea
VM3			Singapore	Singapore
VM4			Sydney	Australia
VM5			Montreal	Canada
VM6			Oregon	USA
VM7	Amazon AWS	64512	Dublin	Ireland
VM8			Virginia	USA
VM9			Mumbai	India
VM10			London	UK
VM11			Frankfurt	Germany
VM12			Paris	France
VM13			Stockholm	Sweden
VM14	Digital Ocean	14061	San Francisco	USA
VM15			Iowa	USA
VM16			Delhi	India
VM17			Tel Aviv	Israel
VM18			Melbourne	Australia
VM19			Johannesburg	South Africa
VM20			Sao Paulo	Brazil
VM21			Hamina	Finland
VM22			Salt Lake City	USA
VM23			Milan	Italy
VM24			Zurich	Switzerland
VM25	Google Cloud	16550	Turin	Italy
VM26			Berlin	Germany
VM27			Mons	Belgium
VM28			Warsaw	Poland
VM29			Doha	Qatar
VM30			Columbus	USA
VM31			Jakarta	Indonesia
VM32			Hong Kong	China
VM33			Taiwan	China
VM34			Santiago	Chile
VM35			Osaka	Japan
VM36			Amsterdam	Netherlands
VM37			Madrid	Spain
VM38			Manchester	United Kingdom
VM39			New York	USA
VM40			Atlanta	USA
VM41			Chicago	USA
VM42			Dallas	USA
VM43	Vultr	20473	Honolulu	USA
VM44			Los Angeles	USA
VM45			Miami	USA
VM46			Seattle	USA
VM47			Silicon Valley	USA
VM48			Mexico City	Mexico
VM49			Toronto	Canada
VM50			Bangalore	India

B Targeted ASes

Table 5 lists the ASes we target in our measurement campaign. Each AS in the Table has a unique identifier with a color code: ASes in red are supposed to have deployed Sr-Mpls according to Cisco. ASes in blue have deployed Sr-Mpls and have responded to our survey (see Sec. 3). ASes in green are confirmed by both Cisco and the survey. The rest (in black) has been randomly selected in the Caida customer cone dataset [17]. We do not know whether they have deployed Sr-Mpls. We classify the ASes according to their position in the AS hierarchy, relying on Caida AS relationship dataset [18]. Among our list, 20% are labeled as Stub, 22% as Content, 45% as Transit and, finally, 13% as Tier-1. The sizes of the initial target lists output by Anaximander (see Sec. 5) are included in the table. They correspond exactly to the number of `traceroute` sent by each VP and intended to hit the targeted AS.

Table 5: List of targeted ASes in our measurement campaign. For each AS, we report the number of traces sent, the number of distinct Ipv4 addresses discovered by our methodology, and whether Sr-Mpls deployment was reported by Cisco and/or through our survey.

AS#ID	ASN	Name	Type	Traces Sent	IPs discovered	SR-MPLS Cisco	SR-MPLS Survey
#1	46467	Dish Network	Stub	2	1	✓	✗
#2	29447	Iliad Italy	Stub	5,888	166	✓	✗
#3	9605	NTT Docomo	Stub	10,034	245	✓	✗
#4	63802	Flets	Stub	512	4	✓	✗
#5	2506	NTT West	Stub	837	18	✓	✗
#6	654	OVH	Stub	0	0	✗	✗
#7	5432	Proximus	Stub	15,392	677	✗	✗
#8	400843	Audacy	Stub	1	0	✗	✗
#9	400322	NGtTel	Stub	15	0	✗	✗
#10	399827	2pifi	Stub	12	4	✗	✗
#11	398872	Big WiFi	Stub	6	2	✗	✗
#12	8835	Binkbroadband	Stub	0	0	✗	✓
#13	45102	Alibaba	Content	14,520	1,813	✓	✗
#14	15169	Google	Content	35,262	19,427	✓	✗
#15	8075	Microsoft	Content	256,419	6,365	✓	✗
#16	138384	Rakuten	Content	1,659	154	✓	✗
#17	17676	Softbank	Content	147,605	21,873	✓	✗
#18	30149	Goldman Sachs	Content	19	10	✗	✗
#19	16509	Amazon	Content	635,599	25,520	✗	✗
#20	14061	Digital Ocean	Content	11,743	3,579	✗	✗
#21	5667	Meta	Content	0	0	✗	✗
#22	43515	YouTube	Content	120	65	✗	✗
#23	138699	Tiktok	Content	14	28	✗	✗
#24	32787	Akamai	Content	4,274	6,988	✗	✗
#25	13335	Cloudflare	Content	10,494	32,735	✗	✗
#26	12322	Free	Transit	42,964	2,024	✓	✗
#27	5410	Bouygues	Transit	27,771	1,048	✓	✗
#28	577	Bell Canada	Transit	29,832	3,748	✓	✗
#29	23764	China Telecom	Transit	11,115	3,374	✓	✗
#30	8220	Colt	Transit	243,811	7,282	✓	✗
#31	2516	KDDI	Transit	89,365	12,994	✓	✗
#32	38631	Line	Transit	423	12	✓	✗
#33	64049	Reliance Jio	Transit	7,014	2,905	✓	✗
#34	132203	Tencent	Transit	7,943	2,922	✓	✗
#35	7018	AT&T	Transit	649,359	44,929	✗	✗
#36	3257	GTT Comm.	Transit	489,738	234,639	✗	✗
#37	6453	Tata Comm.	Transit	275,874	92,854	✗	✗
#38	6762	Telecom Italia	Transit	290,678	32,313	✗	✗
#39	7473	Singtel	Transit	9,549	5,206	✗	✗
#40	6939	Hurricane El.	Transit	652,399	192,324	✗	✗
#41	9002	RETN	Transit	526,697	27,270	✗	✗
#42	2828	Verizon	Transit	26,030	570	✗	✗
#43	7922	Comcast	Transit	272,360	40,382	✗	✗
#44	11232	Midco-Net	Transit	3,153	1,071	✗	✓
#45	13855	CFU-NET	Transit	143	72	✗	✓
#46	293	ESnet	Transit	277,155	307	✗	✓
#47	31034	Aruba	Transit	1,186	346	✗	✓
#48	31631	Elevate	Transit	73	64	✗	✓
#49	32440	Loni	Transit	401	70	✗	✓
#50	33362	Wiktel	Transit	117	39	✗	✓
#51	44092	Halservice	Transit	140	86	✗	✓
#52	7794	Execulink	Transit	599	141	✗	✓
#53	3320	Deutsche Telekom	Tier-1	370,152	65,995	✓	✓
#54	2914	NTT Comm.	Tier-1	504,001	209,589	✓	✗
#55	5511	Orange	Tier-1	51,979	21,376	✓	✗
#56	4637	Telstra	Tier-1	62,075	18,010	✓	✗
#57	1273	Vodafone	Tier-1	24,308	8,248	✓	✗
#58	1299	Arelion	Tier-1	615,851	339,007	✗	✗
#59	174	Cogent	Tier-1	539,127	217,700	✗	✗
#60	3356	Level3	Tier-1	468,812	174,373	✗	✗

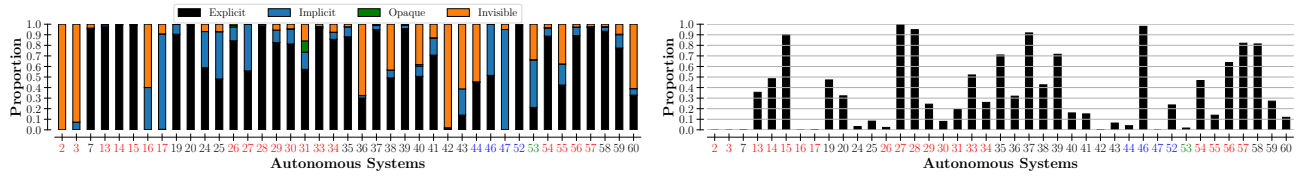

(a) Proportion of explicit, implicit, opaque and invisible MPLS tunnels observed within each AS.

(b) Proportion of paths showing at least one explicit MPLS tunnel.

Figure 13: MPLS tunnels distribution across ASes. AS identifiers match those listed in Table 5. Colors indicate source of SR-MPLS confirmation: red (Cisco), blue (survey), green (both), and black (no explicit confirmation). Identifier ranges reflect AS roles: 1–12 Stub, 13–25 Content Providers, 26–52 Transit, 53–60 Tier-1.

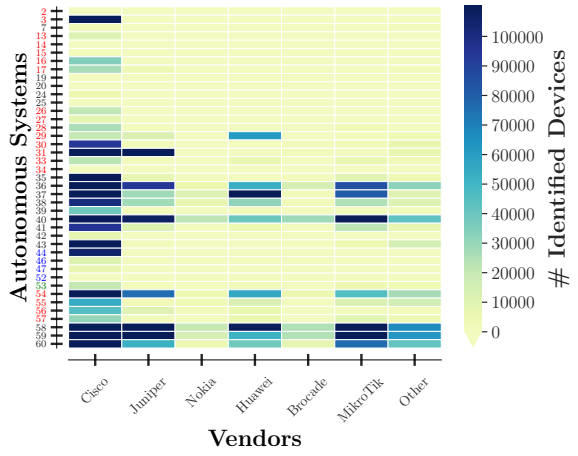

Figure 15: Distribution of identified router vendors per AS using SNMPv3-based fingerprinting. Heatmap's color intensity represents the number of devices attributed to a specific vendor. AS identifiers match those listed in Table 5. Colors indicate source of SR-MPLS confirmation: red (Cisco), blue (survey), green (both), and black (no explicit confirmation). Identifier ranges reflect AS roles: 1–12 Stub, 13–25 Content Providers, 26–52 Transit, 53–60 Tier-1.

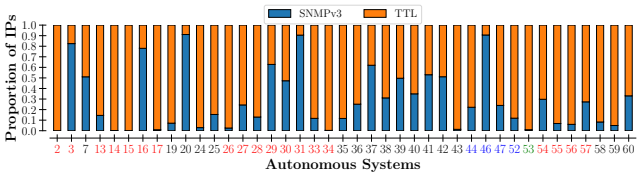

Figure 14: Proportion of SNMPv3-based vs TTL-based fingerprinting across ASes. AS identifiers match those listed in Table 5. Colors indicate source of SR-MPLS confirmation: red (Cisco), blue (survey), green (both), and black (no explicit confirmation). Identifier ranges reflect AS roles: 1–12 Stub, 13–25 Content Providers, 26–52 Transit, 53–60 Tier-1.

C Preliminary Dataset Analysis

Only explicit and, possibly, opaque MPLS tunnels provide conditions to trigger the SR-MPLS detection flags. This is as opposed to implicit and invisible MPLS tunnels that do not quote the LSE in the ICMP `time-exceeded` message.

Explicit tunnels are particularly valuable because they fully expose the MPLS LSE stack for each hop in the `traceroute`, making them eligible for all the SR-MPLS-specific flags defined in the AReST methodology (see Sec. 4). Opaque tunnels, on the other hand, reveal only the last MPLS hop with the LSE, internal MPLS nodes being hidden to `traceroute`, which means they can trigger only the AReST flags that do not rely on a 20-bit label sequence, i.e., LVR, LSVR, and LSO.

Fig. 13a shows how the MPLS tunnels are spread into the four categories in our dataset. The amount of explicit tunnels exceeds the other categories. Opaque tunnels are less frequent while Stub ASes are almost entirely covered by invisible and implicit tunnels, thus preventing any SR-MPLS detection. Furthermore, Fig. 13b shows that a few ASes, that either answered our survey or confirmed their SR-MPLS deployment through Cisco, exhibit a small fraction of explicit tunnels. This is particularly the case for ASes 44 (Midco-Net), 47 (Aruba), and 53 (Deutsche Telekom), among others. These observations reinforce an important point: the SR-MPLS areas uncovered by AReST represent only a lower bound of actual deployment. Many configurations may remain undetectable due to tunnel invisibility.

To assign flags relying on SR-MPLS vendor ranges, we fingerprinted each replying IPv4 interface replying to `traceroute` probing. Overall, we were able to assign a router vendor to approximately 45% of all observed hops in the dataset. Fig. 14 shows that the majority of these mappings, i.e. 88%, came from the TTL-based fingerprinting technique built into TNT [68], the remaining 12% of identified devices were resolved via SNMPv3-based fingerprinting [3].

The SNMPv3-based fingerprinting revealed the presence of several major router vendors in our dataset (see Fig. 15). Cisco devices were the most common by far, followed by Juniper and Huawei. Besides these, we found smaller contributions from other vendors such as Nokia (Alcatel-Lucent) and even a few Linux-based routing platforms. We note that Arista equipment was absent from our results. This is because the public SNMPv3 dataset we used did not contain fingerprints for Arista devices, preventing their identification in our study.

We next examined the distribution of MPLS 20-bit label values observed across all ASes. The 20-bit MPLS label space ranges within

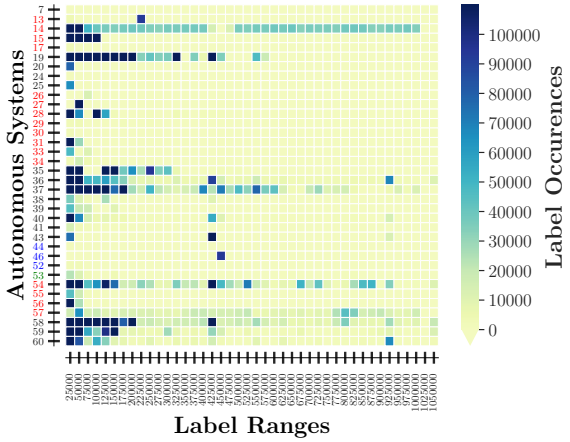

Figure 16: MPLS label range occurrences across ASes. Heatmap's color intensity reflects how often a label in a given range was observed. AS identifiers match those listed in Table 5. Colors indicate source of SR-MPLS confirmation: red (Cisco), blue (survey), green (both), and black (no explicit confirmation). Identifier ranges reflect AS roles: 1–12 Stub, 13–25 Content Providers, 26–52 Transit, 53–60 Tier-1.

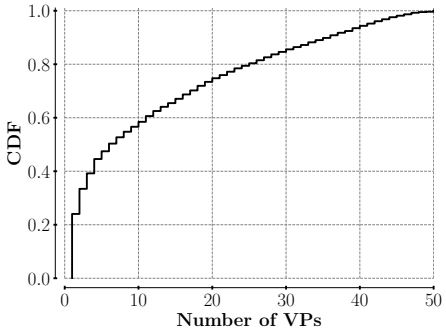

Figure 17: CDF of unique hops discovered as vantage points are added.

$[0; 2^{20}\text{-}1]$, but our dataset's labels were heavily skewed toward low values in that range. Indeed, Fig 16 shows that most MPLS 20-bit labels encountered were relatively small numbers, often in the tens of thousands or less, with very few instances above 100,000. Since most vendors use SR-MPLS label blocks that also reside in a lower portion of the label space, a randomly chosen label from our observations is statistically more likely to fall within an SR-MPLS range than it would be if labels were uniformly spread. In other words, considering that the device vendor has been correctly identified, the skew toward lower label values inherently boosts the probability that some of those labels lie inside known SR-MPLS label ranges (see Table 1 for common SRGB/SRLB values).

Finally, we assess the contribution of each VP to the discovery of new hops and overall path diversity in our dataset. Our measurement used 50 distributed VPs (see Appendix A), and each VP probed the same set of targets from a different perspective. Fig 17 shows a cumulative distribution of unique IPv4 addresses seen as we add more VPs. The observed slow growth indicates that the first few VPs discover a core set of new hops, and subsequent VPs add some additional unique hops, gradually increasing coverage toward 100%. There was no extreme skew where one VP found the majority of hops; instead, the discovery was reasonably well spread out with 50 VPs.

D Ethical Considerations

Probing through `traceroute` and pings may raise an alarm within some operators and, consequently, appear as an attack. To avoid this, we set up our measurement campaign to avoid burdening the network as much as possible. By default, ANAXIMANDER builds the targets list by limiting probing redundancy. Further, we sent probes at a reasonable pace. As such, we believe this work raises no ethical concerns.