Independent market research and competitive analysis of next-generation business and technology solutions for service providers and vendors

HEAVY READING

WHITE PAPER

Will Segment Routing Deliver on Its Promise?

A Heavy Reading white paper produced for Juniper Networks

AUTHOR: STERLING PERRIN, PRINCIPAL ANALYST, HEAVY READING

INTRODUCTION

Segment routing (SR) was introduced to the market just a few years ago with a promise that it would greatly simplify the control plane in IP networks. It promised to move state information from the network to the packet header and provide an end-to-end transport architecture that spans multiple network domains.

Service provider support of SR is growing, but work remains in advance of broad industry adoption. In this white paper, Heavy Reading explores the state of SR today: What is working and what's not? What is the industry doing to ensure broad, reliable adoption at scale for this promising new technology? Topics addressed include the following:

- The current state of SR in service provider networks
- Challenges to broader adoption
- Future use case cases
- Proposed path forward for SR

STATE OF SEGMENT ROUTING TODAY, CURRENT ADOPTION, USE CASES

"Segment routing" and "source routing" are two terms that describe a routing technique in which the source (or ingress) router specifies the route that the packet will take through the network, rather than the path being chosen based on the packet's destination only. SR was officially introduced within the Internet Engineering Task Force (IETF) in 2013 when the Source Packet Routing in Networking (SPRING) Working Group was formed.

SR's primary differentiation from legacy traffic engineering protocols – such as RSVP-TE, LDP, and MPLS-TP[*] – is in its implementation and operational simplicity. While RSVP-TE requires all routers in the network to hold per-path state information, SR requires state only in the source (ingress) router. Removing the requirement to hold state throughout the network leads directly to increased network scale and service agility. Coupling SR with centralized software-defined networking (SDN) control further simplifies some SR uses and is a requirement for others (as discussed later in this paper).

As defined within the IETF, SR is intended to work with all IP forwarding planes. While multiple SR standards are being developed, the SR feature set delivered will be the same regardless of forwarding plane.

[*] RSVP-TE: Resource Reservation Protocol - Traffic Engineering
LDP: Label Distribution Protocol
MPLS-TP: Multiprotocol Label Switching - Transport Profile

Figure 1 shows the different SR standard options supporting the different IP forwarding planes.

Figure 1: Segment Routing Standards and Forwarding Planes

SR Standard	Forward Planes Supported	Status
SR-MPLS	IPv4, MPLS, IPv6 (SRo6)	Complete
SRv6	IPv6	In progress
SRm6	IPv6	In progress

Source: Heavy Reading

Today, there are dozens of IETF drafts related to SR spanning SPRING and, for IPv6 development, the IETF's IPv6 maintenance working group called 6man. SR-Multiprotocol Label Switching (SR-MPLS) standardization is mature and SR-MPLS is running on production networks today. Standardization for the two IPv6-specific implementations – SRv6 and SRm6 – is ongoing. Early SR use cases illustrate the service provider benefits of using the technology.

Traffic Engineering/Traffic Steering

Traditionally, service providers have used RSVP-TE for traffic engineering, but RSVP-TE can be cumbersome and complex to implement. While some operators have deployed RSVP-TE extensively in their networks, others have not. SR offers an alternative to RSVP-TE for traffic engineering and traffic steering use cases.

Even at this early stage, service providers have been able to implement SR for applications in which RSVP-TE has proven too complex to try. Vodafone Germany, for example, has used SR to engineer paths based on latency and application needs and has realized a 50% latency reduction in paths. Microsoft has implemented SR for data center interconnect (DCI) to define different service-level agreements (SLAs) for different applications.

Service providers are just scratching the surface of what can be done. Liquid Telecom has used SR traffic engineering to develop an interesting external-facing application. The service provider has developed a web portal through which customers can set their own constraints to create paths through their networks – putting control back into the hands of the customer, as the operator states. Control in the hands of the customer is a common mantra of operators that have capitalized on the benefits of SDN.

As a final important point, for service providers that have existing RSVP-TE networks, the use of a centralized controller permits the coexistence of RSVP-TE and SR. It also enables the smooth evolution from RSVP-TE to SR over time.

Network Restoration Including Headend Restoration and TI-LFA

Another early-stage use case for SR is network restoration, including headend restoration and Topology Independent – Loop-Free Alternates (TI-LFA). Headend restoration provides end-to-end failure recovery for the entire path that is carrying a TE tunnel's traffic. This restoration is achieved by configuring two paths for connectivity from an SR ingress node to an SR egress node. One path provides primary connectivity while the other path provides backup connectivity. When the SR ingress node detects failure on the primary path, it sends traffic to the backup path. In designating the paths, ensuring segment diversity is critical. Additionally, TI-LFA is a fast-reroute (FRR) mechanism that uses SR to ensure network coverage for 100% of link and node failure scenarios.

Using SR, a backup path is defined for each node and link with a set of labels that route packets around the failure and to the destination. Labels direct packets through all scenarios that would otherwise route them back to failure-affected nodes. A simple label addition mitigates the loop-back challenge through which adjacencies drive packets back to problem sources. Additionally, because the repair tunnel is an SR path, it is not restricted by least cost path rules that prevent 100% coverage guarantees for some predecessor FRR techniques, including LFA and remote LFA (RLFA). Thus, SR TI-LFA is both loop-free and topology independent.

Among the SR use cases, network restoration using either headend restoration or TI-LFA is popular and relatively mature. Significantly, while many of the use cases described in this paper require a controller, it is optional for network restoration. Service providers decide whether they prefer the decisions to be made locally at the node or centrally via a controller.

Stateless Service Function Chaining

Service chaining describes the stringing together of multiple virtual network functions to deliver a specific service to an end customer. For example, an operator may deliver a service consisting of a chain of a load balancer, a firewall, and a proxy. This is an important concept in network functions virtualization (NFV).

SR provides a way to implement stateless service function chaining by programming ordered service chains in segment IDs (SIDs) – either in MPLS label stacks (with SR-MPLS) or directly in IP extension headers (with IPv6). The key benefit of SR in this case is that state does not need to be maintained at each hop in the network. This stateless benefit directly translates to increased simplicity and scalability in networks that implement SR-based service function chaining.

ADOPTION CHALLENGES

Despite early progress, challenges to greater adoption remain. Heavy Reading sees three broad classes of challenges to address, as described below: migration, IPv6 support, and compatibility.

Migration to SDN Control

Some early use cases, such as network restoration with TI-LFA and some simpler traffic steering, work with a traditional distributed control plane. But to realize the full suite of SR use cases, including complex traffic engineering (such as when traffic engineering policy requires bandwidth reservation) and those requiring customer-defined control, centralized controllers are required. At this stage, not all router suppliers offer centralized controllers. Some early adopter service providers have programmed their own controllers from scratch, but only a small subset of service providers have the required software development skills in-house. Mass-market adoption addressing the full SR potential will require commercial controllers available from many sources. Juniper, with its NorthStar controller, is one supplier selling a commercial SDN controller today.

The Road to IPv6 Support

SR-MPLS implementations have achieved some success (as described in early use cases), but the technology still requires an MPLS network, as SIDs are encoded in MPLS labels. Many service providers are happy with SR-MPLS and MPLS, but some look to migrate away from MPLS or – in certain greenfield applications – not use MPLS at all. One of the IETF's goals is to standardize SR for IPv6 networks in which SIDs are encoded directly in IPv6 packets as extension headers. This process eliminates the MPLS requirement from SR.

Yet, the standardization of SR IPv6 implementation is not complete and is currently a matter of debate within the IETF. The primary challenge in SR IPv6 implementation involves the header length. Adding SIDs increases packet length and, in certain use cases, headers will become too long for effective router processing. For example, the use of eight SIDs incurs an overhead of about 20% for an average internet packet size of 500 or 600 bytes.

This challenge has led to the introduction of a compressed routing header (CRH) for SR in IPv6 packets, as proposed in the Segment Routing Mapped to IPv6 (SRm6) IETF draft. By reducing the SID field from 128 bits (as proposed in SRv6) to either 16 or 32 bits, the packet length challenge is largely mitigated. **Figure 2** compares IPv6 header and SID extensions between the two major proposals.

Figure 2: Overhead Comparison between SRv6 and SRm6

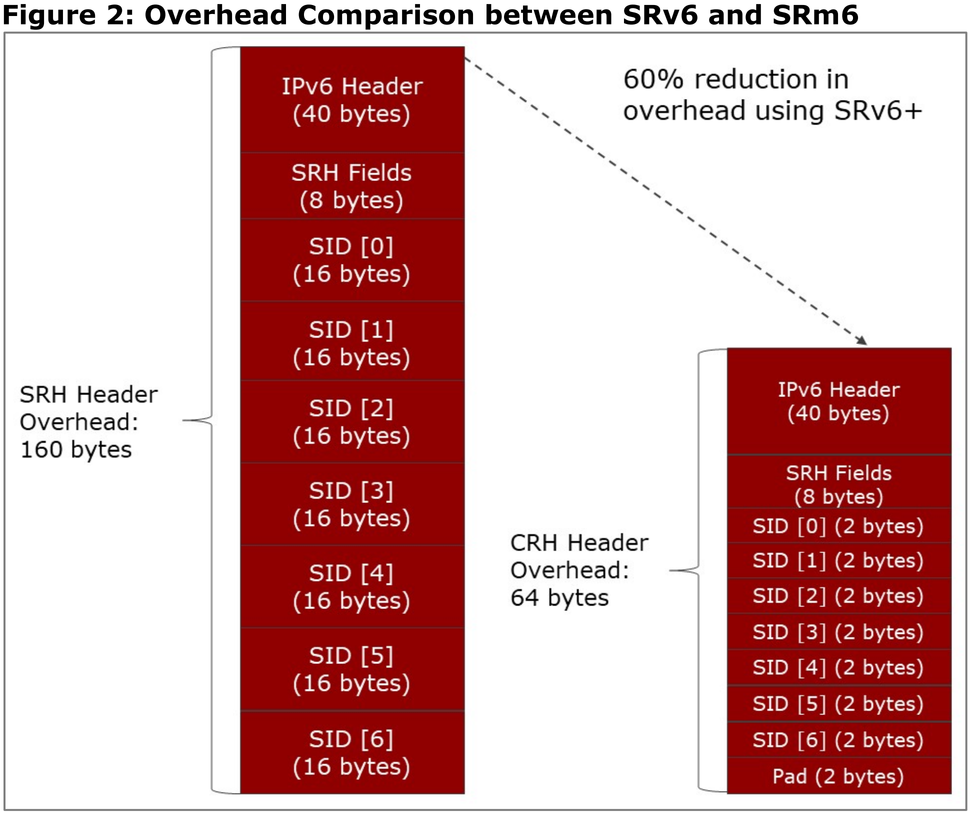

Sources: Heavy Reading, Juniper Networks, 2019

In addition to the SRm6 draft, there are currently five other SRm6-related drafts. Draft contributors include suppliers Juniper and Ericsson and service providers Liquid Telecom, Verizon, NTT Communications, Google, and Baidu.

SR development for IPv6 networks has forked, and two competing approaches have emerged: SRv6 and SRm6. The industry may ultimately need to support two approaches to SR implementation in IPv6 networks.

Compatibility with Installed Base

With any new telecom technology, compatibility with the existing installed base is a chief concern among operators. This is the case with SR as well. Operators will need to maintain MPLS and IPv4 networks for many years, even as they add SR to their networks. Given its use of MPLS, SR-MPLS presents less of a challenge, but RSVP-TE and other protocols must still be accounted for. With SRv6, the challenge becomes harder, as service providers will need to support a new forwarding plane in their IP networks as they transition from IPv4 to IPv6.

FUTURE USE CASES

As SR gains more traction within service providers and as the challenges identified in this paper get resolved, a new set of use cases will open up. In this section, Heavy Reading highlights some of the most compelling emerging use cases on the horizon, but additional high value use cases will emerge over the next couple of years.

5G Network Slicing

End-to-end network slicing will be crucial for operators in offering the full spectrum of 5G use cases, including a mix of Ultra-Reliable Low Latency Communications (URLLC), Massive Machine-Type Communications (mMTC), and Enhanced Mobile Broadband (eMBB). With performance specifications radically different across the spectrum of potential 5G use cases, a single network cannot address them all. At the same time, however, maintaining fully separate networks for different use cases will never be cost-effective. Network slicing allows operators to partition services such that SLAs and key performance indicators (KPIs) – such as latency, throughput, and reliability – can be assured within various "slices" operating within a common network infrastructure.

Within transport, many operators view SR as the most promising option for soft slicing – in which services are partitioned but network resources can still be shared (see **Figure 3** below). Here, the traffic engineering and scale of SR, combined with the centralized control of SDN, are a natural fit. By applying label stacks, operators can quickly define separate slices for an enterprise financial application (which will require the highest levels of reliability and security) and a virtual reality application (which will need lowest latency and lots of bandwidth). Traffic engineering – implemented by the label stacks – goes beyond bandwidth levels and traffic prioritization and can include specific paths through the network and path redundancy.

Figure 3: Are you planning to implement segment routing as part of your 5G transport strategy?

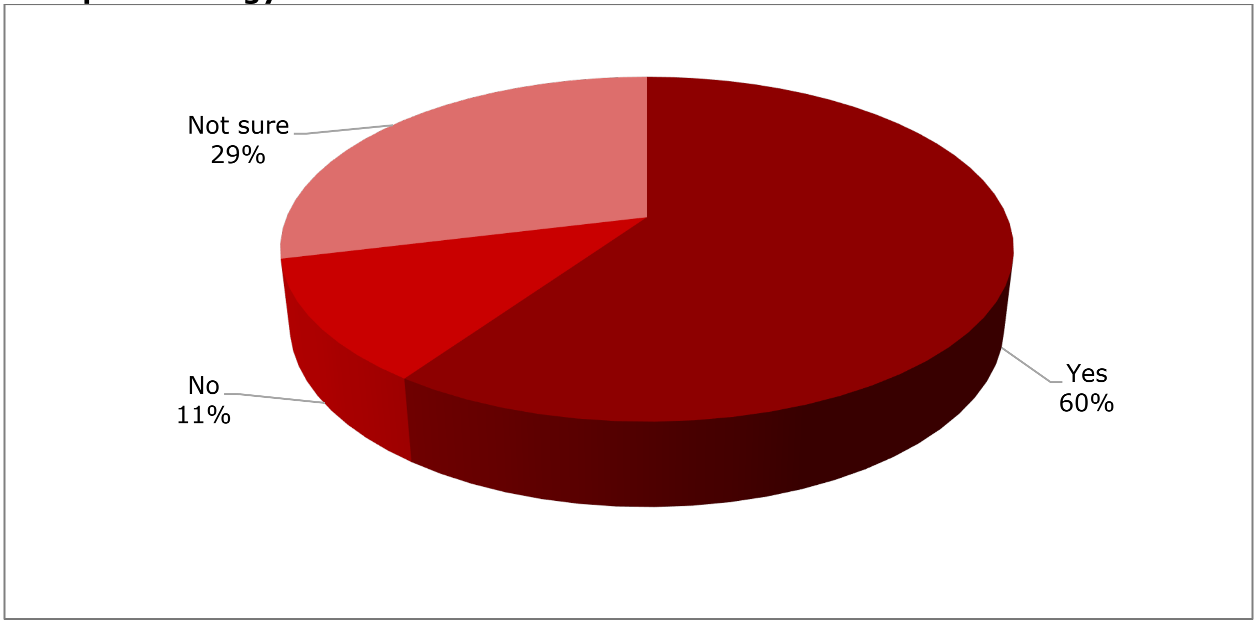

N=104

Source: Operator Strategies for 5G Transport, Sponsored by Ericsson, Fujitsu, Juniper, and Viavi, 2019

Heavy Reading research shows that initial 5G rollouts are focused primarily on the faster broadband speeds (or eMBB use cases), which will not require multiple network slices. Our survey data indicates that URLLC and mMTC use cases are expected to ramp in the 2021-2023 timeframe, following the release of the full 3GPP Release 16 standard in 2020. Support for end-to-end slicing will certainly be required at this time. Still, for end-to-end slicing, standardized coordination between the transport network and the radio access network (RAN) will be needed.

Work to Be Done

The 3GPP is working on standards for slicing and therefore on a de facto reference definition; however, this work is ongoing and is not comprehensive. Within the 3GPP domain, an end-to-end slice can run from the user equipment (e.g., handset) across the RAN and into the core network. While SR can address the transport portion of a 5G network, RAN standardization will be required for the end-to-end use case.

Cross-Domain Stitching to Map Users to Applications on Servers across Multiple Data Centers

Early SR use cases are intra-domain, in which SR steers traffic through a number of routers all residing within the same domain. In these early use cases, if the service provider uses a centralized controller, that controller acts only on the routers that reside within the domain. Thus, segment-routed networks exist as separate islands across service providers, or even within service providers (in cases where there are multiple domains).

One real-world example is data center connectivity. Service providers use SR-MPLS today to steer traffic within their data centers. Exiting the data center, however, traffic typically travels over an MPLS backbone to reach a destination data center. The destination data center may employ SR within its own domain.

Thus, while SR provides benefits for intra-data center or intra-domain traffic steering today, some kind of cross-domain stitching is required in order to steer traffic end-to-end from one domain to another. Today, no universal cross-domain stitching technology exists. However, given its position as a next-gen traffic steering technology within domains, SR is a logical choice to provide this full end-to-end functionality for the future.

Work to Be Done

There are challenges to be addressed, as MPLS was not initially envisioned for this function. One problem is that the size of the MPLS label stacks increase linearly with the length of the source route, and MPLS implementations may not support the full label stack required. Additionally, steering across domains requires coordination between the controllers. If domains are operated by different entities, the controllers may not be able (or willing) to share the required level of information.

Segment Routing and SD-WAN

The software-defined wide-area network (SD-WAN) has emerged as a strong connectivity option for enterprise cloud applications and is poised for strong growth over the next 5 years and beyond. Ovum forecasts global SD-WAN services revenue will increase at a 21% compound annual growth rate (CAGR) from 2018 to 2023 while global enterprise service revenue remains flat.

Although some companies position SR and SD-WAN as competitive technologies, Heavy Reading views them as complementary innovations. When combined, these technologies can provide significant differentiation for operator SD-WAN services. Bell Canada is one operator championing this future use case, based on the complementary nature of SD-WAN and SR.

The ability to set per-flow policy mechanisms is one of the most important aspects of SR for the SD-WAN use case, providing service providers (and their enterprise customers) the ability to coordinate between the overlay (SD-WAN) and the underlay (IP transport) networks in a way that has not been possible. Thus, for the same IP destination, customers can have different paths through the network based on the per-flow policy. For latency bound services, for example, SR can specify latency and diversity for the path taken.

Work to Be Done

SD-WAN and SR interaction will typically require two controllers: the SD-WAN controller for the SD-WAN and the SR controller, which can be based on the Path Computation Element Protocol (PCEP) or NETCONF. The SD-WAN controller may be either owned by the enterprise customer or managed by the operator (in a managed service). Interaction between the SD-WAN and SR controllers will be required for direct coordination. For multi-vendor networks, such interaction will require a standard interface through which the SD-WAN and the SR controllers communicate, but these standards do not yet exist.

CONCLUSIONS AND FUTURE OF SEGMENT ROUTING

SR promises to greatly simplify the operations and management of IP networks. Service provider demand for SR is strong and early deployments have focused on topology independent restoration, traffic steering, traffic engineering, and some others. Emerging use case possibilities are almost endless, including complex traffic engineering applications, SR for 5G network slicing, cross-domain stitching, SR underlay in SD-WANs, and many more.

Still, there is work to be done for SR to realize its true potential. While some use cases are effective in controllerless networks (such as traffic steering), many benefit from centralized control, and some use cases absolutely require the centralized network view (such as traffic engineering with bandwidth reservation). SR with centralized control will become increasingly important as the technology matures.

There is also industry confusion regarding the applicability and future of SR in IPv6 networks. For IPv6 implementation, the industry has reached a fork, with two IETF standards in development: SRv6 and SRm6. The industry must decide whether it will adopt one or both. The good news is that feature sets will be the same regardless of the data plane used. SR-MPLS can address most current and future use cases as long as network equipment supports MPLS.