SPRING Working Group                                             S. Peng
Internet-Draft                                                     Z. Li
Intended status: Informational                       Huawei Technologies
Expires: January 9, 2020                                          C. Xie
                                                                   C. Li
                                                           China Telecom
                                                            July 8, 2019

# SRv6 Compatibility with Legacy Devices
# draft-peng-spring-srv6-compatibility-01

Abstract

When deploying SRv6 on legacy devices, there are some compatibility challenges such as the support of SRH processing.

This document identifies some of the major challenges, and provides solutions that are able to mitigate those challenges and smooth the migration towards SRv6 deployment.

## Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Copyright Notice



Table of Contents



## 1. Introduction

Segment Routing (SR) is a source routing paradigm, which allows a headend node to steer the packets through an ordered list of instructions, i.e. segments [RFC8402]. A segment can either be topological or service based. SR over IPv6 (SRv6) [I-D.filsfils-spring-srv6-network-programming] is the SR instantiated on the IPv6 data plane with a new type of routing extension header, i.e. SR Header (SRH) [I-D.ietf-6man-segment-routing-header]. An SRv6 segment, also called SRv6 SID, is a 128-bit value, represented as

LOC:FUNCT:ARGS (ARGS is optional), and encoded as an IPv6 address. An ordered list of SRv6 SIDs forms an SR Policy, which can be used for, for example, Traffic Engineering (TE), Service Function Chaining (SFC), and In-situ Operations, Administration, and Maintenance (IOAM). Meanwhile, it will also bring challenges on the legacy devices to support SRv6 correspondingly.

This document provides solutions that can mitigate the identified compatibility challenges and ease the evolution towards SRv6 deployment.

## 2. Compatibility challenges

By adopting SR Policy, the states in the network can be greatly reduced, which will relieve the devices and evolve into stateless fabric ultimately. However, it will also bring compatibility challenges on the legacy devices correspondingly. In particular, the legacy devices need to upgrade in order to support the processing of SRH.

Furthermore, as the segments in the segment list increase the SR Policy incrementally expends, the encapsulation header overhead increases, which will also impose high requirements on the performance of hardware forwarding (i.e. the capability of chipset).

This section identifies the imposed challenges in the following SPRING use cases.

### 2.1. Fast Reroute (FRR)

FRR is deployed to cope with link or node failures by precomputing backup paths. By relying on SR, Topology Independent Loop-free Alternate Fast Re-route (TI-LFA) [I-D.bashandy-rtgwg-segment-routing-ti-lfa] provides a local repair mechnism with the ability to activate the data plane switch-over onto a loop free backup path irrespective of topologies prior and after the sudden failure.

Using SR, there is no need to create state in the network in order to enforce FRR behavior. Correspondingly, the Point of Local Repair, i.e. the protecting router, needs to insert a repair list at the head of the segment list in the SR header, encoding the explicit post-convergence path to the destination. This action will increase the length of the segment list in the SRH as shown in Figure 1.

## 2.2. Traffic Engineering (TE)

TE enables operators to control specific traffic flows going through configured explicit paths. There are loose and strict options. With the loose option, only a small number of hops along the paths are explicitly expressed, while the strict option specifies each individual hop in the explicit path, e.g. to encode a low-latency path from node A to node B.

With SRv6, the strict source-routed explicit paths will result in a long segment list in the SRH as shown in Figure 1, which places high requirements on the devices.

## 2.3. Service Function Chaining (SFC)

The SR segments can also encode instructions, called service segments, for steering packets through services running on physical service appliances or virtual network functions (VNF) running in a virtual enviornment [I-D.xuclad-spring-sr-service-programming]. These service segments can also be integrated in an SR policy along with node and adjacency segments. This feature of SR will further increase the length of the segment list in the SRH as shown in Figure 1.

In terms of SR awareness, there are two types of services, i.e. SR-aware and SR-unaware services, which both impose new requirements on the hardware. The SR-aware service needs to be fully capable of processing SR traffic, while for the SR-unaware services, an SR proxy function needs to be defined.

If the Network Service Header (NSH) based SFC [RFC8300] has already been deployed in the network, the compatibility with existing NSH is required.

## 2.4. IOAM

IOAM, i.e. "in-situ" Operations, Administration, and Maintenance (OAM), encodes telemetry and operational information within the data packets to complement other "out-of-band" OAM mechnisms, e.g. ICMP and active probing. The IOAM data fields, i.e. a node data list, hold the information collected as the packets traversing the IOAM domain [I-D.ietf-ippm-ioam-data], which is populated iteratively starting with the last entry of the list.

The IOAM data can be embedded into a variety of transports. To support the IOAM on the SRv6 data plane, the O-flag in the SRH is defined [I-D.ali-spring-srv6-oam], which implements the "punt a timestamped copy and forward" or "forward and punt a timestamped

copy" behavior. The IOAM data fields, i.e. the node data list, are encapsulated in the IOAM TLV in SRH, which further increases the length of the SRH as shown in Figure 1.

```
                                                            +-----------+
                                                            |IPv6 packet|
                                                            +-----------+
                                                            /           /
                                             +-----------+  / IOAM Info /
                                             |IPv6 packet|  /           /
                              +-----------+  +-----------+  +-----------+
                              |IPv6 packet|  /           /  /           /
               +-----------+  +-----------+  /           /  /           /
               |IPv6 packet|  /           /  / SF Chain  /  / SF Chain  /
+-----------+  +-----------+  /  TE Path  /  /           /  /           /
|IPv6 packet|  /TI-LFA Path/  /           /  /           /  /           /
+-----------+  +-----------+  +-----------+  +-----------+  +-----------+
|SA,DA      |  |SA,DA      |  |SA,DA      |  |SA,DA      |  |SA,DA      |
+-----------+  +-----------+  +-----------+  +-----------+  +-----------+
   SRv6 BE       SRv6 BE+        SRv6 TE       SRv6 SFC       SRv6 SFC+
                 TI-LFA                                          IOAM
```

Figure 1. Evolution of SRv6 SRH

The compatibility challenges on the legacy devices are summarised as follows,

- o The legacy devices need to upgrade in order to support the processing of SRH
- o As the SRH expands, the overhead increases and correspondingly the effective payload decreases
- o As the SRH expands, the hardware forwarding performance reduces which requires high capability of chipset

## 3. Solutions

This section provides solutions to mitigate the above-mentioned challenges.

### 3.1. TE

With the strict traffic engineering, the resulted long SID list in the SRH raises high requirements on the hardware chipset, which can be mitigated by the following solutions.

### 3.1.1. Binding SID (BSID)

Binding SID involves a list of SIDs, and is bound to an SR Policy. The node(s) that imposes the bound policy needs to store the SID list. When a node receives a packet with its active segment as a BSID, the node will steer the packet onto the bound policy accordingly.

To reduce the long SID list of a strict TE explicit path, BSID can be used at the selected nodes, maybe according to the processing capacity of the hardware chipset. BSID can also be used to impose the repair list in the TI-LFA as described in Section 2.1.

### 3.1.2. PCEP FlowSpec

When the SR architecture adopts a centralized model, the SDN controller (e.g. Path Computation Element (PCE)) only needs to apply the SR policy at the head-end. There is no state maintained at midpoints and tail-ends. Eliminating states in the network (midpoints and tail-points) is a key benefit of utilizing SR. However, it also leads to a long SID list for expressing a strict TE path.

PCEP FlowSpec [I-D.ietf-pce-pcep-flowspec] provides a trade-off solution. PCEP FlowSpec is that PCEP with a set of extensions is able to disseminate Flow Specifications (i.e. filters and actions) to allow indicating how the classified traffic flows will be treated. In an SR-enabled network, PCEP FlowSpec can be applied at the midpoints to enforce traffic engineering policies where it is needed. In that case, states need to be maintained at the corresponding midpoints of a TE explicit path, but the SID list can be shortened.

## 3.2. SFC

Currently two approaches are proposed to support SFC over SRv6, i.e. stateless SFC [I-D.xuclad-spring-sr-service-programming] and stateful SFC [I-D.guichard-spring-nsh-sr].

### 3.2.1. Stateless SFC

A service can also be assigned an SRv6 SID which is integrated into an SR policy and used to steer traffic to it. In terms of the capability of processing the SR information in the received packets, there are two types of services, i.e. SR-aware service and SR-unware service. An SR-aware service is capable of processing the SRH in the received packets. While an SR-unaware service, i.e. legacy service, is not able to process the SR information in the traffic it receives, and may drop the received packets. In order to support such services

in an SRv6 domain, the SR proxy is introduced to handle the processing of SRH on behalf of the SR-unware service. The service SID associated with the SR-unaware service is instantiated on the SR proxy, which is used to steer traffic to the service.

The SR proxy intercepts the SR traffic destined for the service via the locally instantiated service SID, removes the SR information, and sends the non-SR traffic out on a given interface to the service. When receiving the traffic coming back from the service, the SR proxy will restore the SR information and forwards it to the next segment in the segment list.

### 3.2.2. Stateful SFC

The NSH and SR can actually be integrated in order to support SFC in an efficient and cost-effective manner while maintaining separation of the service and transport planes .

In this NSH-SR integration solution, NSH and SR work jointly and complement each other. Specifically, SR is responsible for steering packets along a given Service Function Path (SFP) while NSH is for maintaining the SFC instance context, i.e. Service Path Identifier (SPI), Service Index (SI), and any associated metadata.

When a service chain is established, a packet associated with that chain will be first encapsulated with an NSH and then an SRH, and forwarded in the SR domain. When the packet arrives at an SFF and needs to be forwarded to an SF, the SFF performs a lookup based on the service SID associated with the SF to retrieve the next-hop context (a MAC address) between the SFF and SF. Then the SFF strips the SRH and forwards the packet with NSH carrying metadata to the SF where the packet will be processed as specified in [RFC8300]. In this case, the SF is not required to be capable of the SR operation, neither is the SR proxy. Meanwhile, the stripped SRH will be updated and stored in a cache in the SFF, indexed by the NSH SPI for the forwarding of the packet coming back from the SF.

## 3.3. Light Weight IOAM

In most cases, after the IPv6 Destination Address (DA) is updated according to the active segment in the SRH, the SID in the SRH will not be used again. However, the entire SID list in the SRH will still be carried in the packet along the path till a PSP/USP is enforced.

The light weight IOAM method [I-D.li-spring-passive-pm-for-srv6-np] makes use of the used segments in the SRH to carry the IOAM

information, which saves the extra space in the SRH and mitigate the requirements on the hardware.

## 4. Summary

The SRH enables a great number of features for SRv6 and opens new network programming possibilities. By using SRH, it relieves the network devices from states, evolving towards stateless fabric, while the complexity in the control plane increases. The corresponding challenges imposed on the hardware chipset become high as the SRH expands when supporting the diverse use cases. The trade-off solutions presented in this document are able to mitigate these challenges and smooth the evolution in operators' networks.

## 5. IANA Considerations

There are no IANA considerations in this document.

Note to RFC Editor: this section may be removed on publication as an RFC.

## 6. Security Considerations

TBD

## 7. Acknowledgements

TBD

## 8. Normative References

[I-D.ali-spring-srv6-oam]
Ali, Z., Filsfils, C., Kumar, N., Pignataro, C., faiqbal@cisco.com, f., Gandhi, R., Leddy, J., Matsushima, S., Raszuk, R., daniel.voyer@bell.ca, d., Dawra, G., Peirens, B., Chen, M., and G. Naik, "Operations, Administration, and Maintenance (OAM) in Segment Routing Networks with IPv6 Data plane (SRv6)", draft-ali-spring-srv6-oam-02 (work in progress), October 2018.

[I-D.bashandy-rtgwg-segment-routing-ti-lfa]
Bashandy, A., Filsfils, C., Decraene, B., Litkowski, S., Francois, P., daniel.voyer@bell.ca, d., Clad, F., and P. Camarillo, "Topology Independent Fast Reroute using Segment Routing", draft-bashandy-rtgwg-segment-routing-ti-lfa-05 (work in progress), October 2018.

[I-D.filsfils-spring-srv6-network-programming]
Filsfils, C., Camarillo, P., Leddy, J., daniel.voyer@bell.ca, d., Matsushima, S., and Z. Li, "SRv6 Network Programming", draft-filsfils-spring-srv6-network-programming-07 (work in progress), February 2019.

[I-D.guichard-spring-nsh-sr]
Guichard, J., Song, H., Tantsura, J., Halpern, J., Henderickx, W., Boucadair, M., and S. Hassan, "NSH and Segment Routing Integration for Service Function Chaining (SFC)", draft-guichard-spring-nsh-sr-01 (work in progress), March 2019.

[I-D.ietf-6man-segment-routing-header]
Filsfils, C., Dukes, D., Previdi, S., Leddy, J., Matsushima, S., and d. daniel.voyer@bell.ca, "IPv6 Segment Routing Header (SRH)", draft-ietf-6man-segment-routing-header-21 (work in progress), June 2019.

[I-D.ietf-ippm-ioam-data]
Brockners, F., Bhandari, S., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Mozes, D., Lapukhov, P., Chang, R., daniel.bernier@bell.ca, d., and J. Lemon, "Data Fields for In-situ OAM", draft-ietf-ippm-ioam-data-06 (work in progress), July 2019.

[I-D.ietf-pce-pcep-flowspec]
Dhody, D., Farrel, A., and Z. Li, "PCEP Extension for Flow Specification", draft-ietf-pce-pcep-flowspec-03 (work in progress), February 2019.

[I-D.li-spring-passive-pm-for-srv6-np]
Li, C. and M. Chen, "Passive Performance Measurement for SRv6 Network Programming", draft-li-spring-passive-pm-for-srv6-np-00 (work in progress), March 2018.

[I-D.xuclad-spring-sr-service-programming]
Clad, F., Xu, X., Filsfils, C., daniel.bernier@bell.ca, d., Li, C., Decraene, B., Ma, S., Yadlapalli, C., Henderickx, W., and S. Salsano, "Service Programming with Segment Routing", draft-xuclad-spring-sr-service-programming-02 (work in progress), April 2019.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8300] Quinn, P., Ed., Elzur, U., Ed., and C. Pignataro, Ed., "Network Service Header (NSH)", RFC 8300, DOI 10.17487/RFC8300, January 2018, <https://www.rfc-editor.org/info/rfc8300>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <https://www.rfc-editor.org/info/rfc8402>.

Authors' Addresses

Shuping Peng
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: pengshuping@huawei.com

Zhenbin Li
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: lizhenbin@huawei.com

Chongfeng Xie
China Telecom
China Telecom Information Science&Technology Innovation park, Beiqijia
Town,Changping District
Beijing 102209
China

Phone: +86-10-50902116
Email: xiechf.bri@chinatelecom.cn

Cong Li
China Telecom
China Telecom Information Science&Technology Innovation park, Beiqijia
Town,Changping District
Beijing 102209
China

Phone: +86-10-50902556

Email: licong.bri@chinatelecom.cn