Network Working Group
Internet-Draft
Intended status: Informational
Expires: December 31, 2014

X. Xu
Z. Li
Huawei
H. Shah
Ciena
L. Contreras
Telefonica I+D

June 29, 2014

Service Function Chaining Use Case for SPRING
draft-xu-spring-sfc-use-case-02

Abstract

This document describes a particular use case for SPRING where the Segment Routing mechanism is leveraged to realize the service path layer functionality of the Service Function Chaining (i.e, steering traffic through the service function path).

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

When applying a particular Service Function Chaining (SFC) [I-D.quinn-sfc-arch] to the traffic selected by the service classifier, the traffic need to be steered through an ordered set of service nodes in the network. This ordered set of service nodes indicates the service function path which is actually the instantiation of the above SFC in the network. Furthermore, additional information about the traffic (a.k.a. metadata) which is helpful for enabling value-added services may need to be carried across those service nodes within the SFC instantiation. As mentioned in [I-D.rijsman-sfc-metadata-considerations] "...it is important to make a distinction between fields which are used at the service path layer to identify the Service Path Segment, and additional fields which carry metadata which is imposed and interpreted at the service function layer. Combining both types of fields into a single header should probably be avoided from a layering point of view. "

Segment Routing (SR) [I-D.filsfils-spring-segment-routing] is a source routing paradigm which can be used to steer traffic through an ordered set of routers. SR can be applied to the MPLS data plane [I-D.gredler-spring-mpls] [I-D.filsfils-spring-segment-routing-mpls]and the IPv6 data plane [I-D.previdi-6man-segment-routing-header].

This document describes a particular use case for SPRING where the SR mechanism is leveraged to realize the service path layer

functionality of the SFC (i.e, steering traffic through the service function path).

1.1. 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].

2. Terminology

This memo makes use of the terms defined in [I-D.filsfils-spring-segment-routing] and [I-D.quinn-sfc-arch].

3. SFC Use Case

```
           +----------------------------------------------- ----+
           |                    SR Netowrks                     |
           |            +---------+       +---------+           |
           |            | +-----+ |       | +-----+ |           |
           |       (1)  | | SF1 | |  (2)  | | SF2 | | (3)       |
      +----+-----+ ---> | +-----+ | ----> | +-----+ | --->  +---+---+
      |Classifier+------+   SN1   +-------+   SN2   +-------+   D   |
      +----------+      +---------+       +---------+       +---+---+
           |                                                    |
           +----------------------------------------------------+
            Figure 1: Service Function Chaining in SR Networks
```

As shown in Figure 1, assume SN1 and SN2 are two SR-capable nodes meanwhile they are service nodes which offer service function SF1 and SF2 respectively. In addition, they have allocated and advertised segment IDs (SID) for the service functions they are offering. For example, SN1 allocates and advertises an SID, i.e., SID(SF1) for service function SF1 while SN2 allocates and advertises an SID, i.e., SID(SF2) for service function SF2. These SIDs which are used to indicate service functions are referred to as Service Function SIDs. In addition, assume the node SIDs for SN1 and SN2 are SID(SN1) and SID(SN2) respectively.

How to steer a packet through a service fucntion path in both MPLS-SR and IPv6-SR cases is illustrated in the following two sub-sections respectively.

3.1. SFC in MPLS-SR Case

In the MPLS-SR case, those service function SIDs as mentioned above would be interpreted as local MPLS labels. Meanwhile, to simplify

the illustrationIn in this document, those node SIDs as mentioned above would be interpreted as MPLS global labels.

Now assume a given packet destined for destination D is required to go through a service function chain {SF1, SF2} before reaching its final destination D. The service classifier therefore would attach a segment list {SID(SN1), SID(SF1), SID(SN2), SID(SF2)} to the packet. This segment list is actually represented by a MPLS label stack. In addition, the service classifier could optionally impose metadata on the packet through the Network Service Header (NSH) [I-D.quinn-sfc-nsh]. Here the Service Path field wihin the NSH would not be used for the path selection purpose anymore and therefore it MUST be set to a particular value to indicate such particular usage. In addition, the service index value within the NSH is set to 2 since there are two service nodes within the service function path. How to impose the NSH on a MPLS packet is outside the scope of this document. When the encapsulated packet arrives at SN1, SN1 would know which service function should be performed according to SID (SF1). If a NSH is carried in that packet, SN1 could further consume the metadata contained in the NSH and meanwhile decrease the service index value within the NSH by one. When the encapsualted packet arrives at SN2, SN2 would do the similar action as what has been done by SN1. Furthermore, since SN2 is the last service node within the service function path, SN2 MUST strip the NSH (if it has been imposed) before sending the packet to D.

3.2. SFC in IPv6-SR Case

In the IPv6-SR case, those service function SIDs as mentioned above would be interpreted as IPv6 link-local addresses while those node SIDs as mentioned above would be interpreted as IPv6 global unicast addresses.

Now assume a given IPv6 packet destined for destination D is required to go through a service function chain {SF1, SF2} before reaching its final destination D. The service classifier therefore would attach a SR header containing a segment list {SID(SF1), SID(SN2),SID(SF2),SID(D)} to the IPv6 packet. This segment list is actually represented by an ordered list of IPv6 addresses. The IPv6 destination address is filled with SID(SN1). In addition, the service classifier could optionally impose metadata on the above IPv6 packet through the NSH and meanwhile carry the original IPv6 source address in the Original Source Address field of the packet. When the above IPv6 packet arrives at SN1, SN1 would know which service function should be performed according to SID (SF1). If a NSH is carried in that packet, SN1 could further consume the metadata contained in the NSH and meanwhile decrease the service index value within the NSH by one. When the packet arrives at SN2, SN2 would do

the similar action as what has been done by SN1. Furthermore, since SN2 is the second last node in the segment list, SN2 should strip the SR header and meanwhile fill in the IPv6 source address with the Original Source Address (if available) before sending the packet towards D. Besides, since SN2 is the last service node within the service path, SN2 MUST strip the NSH (if it has been imposed) before sending the packet to D.

4. Acknowledgements

TBD.

5. IANA Considerations

TBD.

6. Security Considerations

This document does not introduce any new security risk.

7. References

7.1. Normative References

[I-D.filsfils-spring-segment-routing-mpls]
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Milojevic, I., Shakir, R., Ytti, S., Henderickx, W., Tantsura, J., and E. Crabbe, "Segment Routing with MPLS data plane", draft-filsfils-spring-segment-routing-mpls-02 (work in progress), June 2014.

[I-D.gredler-spring-mpls]
Gredler, H., Rekhter, Y., Jalil, L., Kini, S., and X. Xu, "Supporting Source/Explicitly Routed Tunnels via Stacked LSPs", draft-gredler-spring-mpls-06 (work in progress), May 2014.

[I-D.previdi-6man-segment-routing-header]
Previdi, S., Filsfils, C., Field, B., and I. Leung, "IPv6 Segment Routing Header (SRH)", draft-previdi-6man-segment-routing-header-01 (work in progress), June 2014.

[I-D.quinn-sfc-arch]
Quinn, P. and J. Halpern, "Service Function Chaining (SFC) Architecture", draft-quinn-sfc-arch-05 (work in progress), May 2014.

[I-D.quinn-sfc-nsh]
Quinn, P., Guichard, J., Fernando, R., Surendra, S., Smith, M., Yadav, N., Agarwal, P., Manur, R., Chauhan, A., Elzur, U., McConnell, B., and C. Wright, "Network Service Header", draft-quinn-sfc-nsh-02 (work in progress), February 2014.

[I-D.rijsman-sfc-metadata-considerations]
Rijsman, B. and J. Moisand, "Metadata Considerations", draft-rijsman-sfc-metadata-considerations-00 (work in progress), February 2014.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

7.2. Informative References

[I-D.filsfils-spring-segment-routing]
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Milojevic, I., Shakir, R., Ytti, S., Henderickx, W., Tantsura, J., and E. Crabbe, "Segment Routing Architecture", draft-filsfils-spring-segment-routing-03 (work in progress), June 2014.

Authors' Addresses

Xiaohu Xu
Huawei

Email: xuxiaohu@huawei.com

Zhenbin Li
Huawei

Email: lizhenbin@huawei.com

Himanshu Shah
Ciena

Email: hshah@ciena.com

Luis M. Contreras
Telefonica I+D
Ronda de la Comunicacion, s/n
Sur-3 building, 3rd floor
Madrid, 28050
Spain

Email: luismiguel.contrerasmurillo@telefonica.com
URI: http://people.tid.es/LuisM.Contreras/