

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 9, 2015

X. Xu
Z. Li
Huawei
H. Shah
Ciena
L. Contreras
Telefonica I+D
March 8, 2015

**Service Function Chaining Using MPLS-SPRING
draft-xu-sfc-using-mpls-spring-03**

Abstract

Source Packet Routing in Networking (SPRING) WG specifies a special source routing mechanism. Such source routing mechanism can be leveraged to realize the service path layer functionality of the service function chaining (i.e, steering traffic through a particular service function path) by encoding the service function path or the service function chain information as the explicit path information. This document describes how to leverage the MPLS-based source routing mechanism as developed by the SPRING WG to realize the service path layer functionality of the service function chaining.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 9, 2015.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Requirements Language	3
2.	Terminology	3
3.	Solution Description	3
3.1.	Encoding SFP Information by an MPLS Label Stack	4
3.2.	How to Contain Metadata within an MPLS Packet	5
4.	Acknowledgements	5
5.	IANA Considerations	5
6.	Security Considerations	5
7.	References	5
7.1.	Normative References	5
7.2.	Informative References	6
	Authors' Addresses	6

[1.](#) Introduction

When applying a particular Service Function Chain (SFC) [[I-D.ietf-sfc-architecture](#)] to the traffic selected by a service classifier, the traffic need to be steered through an ordered set of Service Functions (SF) in the network. This ordered set of SFs in the network indicates the Service Function Path (SFP) associated with the above SFC. To steer the selected traffic through an ordered list of SFs in the network, the traffic need to be attached by the service classifier with the information about the SFP (i.e., specifying exactly which Service Function Forwarders (SFFs) and which SFs are to be visited by traffic), the SFC, or the partially specified SFP which is in between the former two extremes. Source Packet Routing in Networking (SPRING) WG specifies a special source routing mechanism which can be used to steer traffic through an ordered set of routers (i.e., an explicit path). Such source routing mechanism can be leveraged to realize the service path layer functionality of the SFC (i.e., steering traffic through a particular SFP) by encoding the SFP information as the explicit path information contained in packets. The source routing mechanism specified by the SPRING WG can be applied to the MPLS data plane [[I-D.ietf-spring-segment-routing-mpls](#)]. This document describes how

to leverage the MPLS-based source routing mechanisms to realize the service path layer functionality of the service function chaining. Note that this approach is aligned with the Transport Derived SFF mode as described in Section 4.3.1 of [[I-D.ietf-sfc-architecture](#)].

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.ietf-spring-segment-routing](#)] and [[I-D.ietf-sfc-architecture](#)].

3. Solution Description

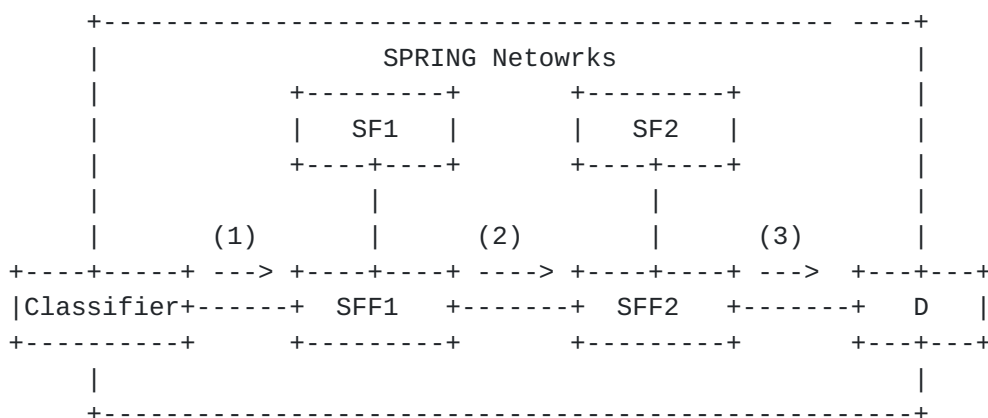


Figure 1: Service Function Chaining in SPRING Networks

As shown in Figure 1, assume SFF1 and SFF2 are two MPLS-SPRING-capable nodes. They are also Service Function Forwarders (SFF) to which two SFs (i.e., SF1 and SF2) are attached respectively. In addition, they have allocated and advertised Segment IDs (SID) for their locally attached SFs. In the MPLS-SPRING context, SIDs are intercepted as MPLS labels. For example, SFF1 allocates and advertises an SID (i.e., SID(SF1)) for SF1 while SFF2 allocates and advertises an SID (i.e., SID(SF2)) for SF2. These SIDs which are used to indicate SFs are referred to as SF SIDs. To encode the SFP information by an MPLS label stack, those SF SIDs as mentioned above would be interpreted as local MPLS labels. In addition, assume node SIDs for SFF1 and SFF2 are SID(SFF1) and SID(SFF2) respectively. Now assume a given traffic flow destined for destination D is selected by the service classifier to go through a particular SFC (i.e., SF1->SF2) before reaching its final destination D. [Section 3.1](#) describes how to leverage the MPLS-based source routing mechanisms to realize

the service path functionality of the service function chaining (i.e., by encoding the SFP information within an MPLS label stack). [Section 3.2](#) describes how to carry metadata over MPLS packets.

3.1. Encoding SFP Information by an MPLS Label Stack

Since the selected packet needs to travel through an SFC (i.e., SF1->SF2), the service classifier would attach a segment list of (i.e., SID(SFF1)->SID(SF1)->SID(SFF2)-> SID(SF2)) which indicates the corresponding SFP to the packet. This segment list is actually represented by a MPLS label stack. When the encapsulated packet arrives at SFF1, SFF1 would know which SF should be performed according to the current top label (i.e., SID (SF1)) of the received MPLS packet. Before sending the packet to SF1, the whole MPLS label stack (i.e., SID(SFF2)->SID(SF2)) MUST be stripped. After receiving the packet returned from SF1, SFF1 would reimpose the MPLS label stack which had been stripped before to the packet and then send it to SFF2 according to the current top label (i.e., SID (SFF2)). When the encapsulated packet arrives at SFF2, SFF2 would do the similar action as what has been done by SFF1. To some extent, the MPLS label stack here could be looked as a specific implementation of the SFC encapsulation used for containing the SFP information [[I-D.ietf-sfc-architecture](#)].

If there is no MPLS LSP towards the next node segment (i.e., the next SFF identified by the current top label), the corresponding IP-based tunnel (e.g., MPLS-in-IP/GRE tunnel [[RFC4023](#)], MPLS-in-L2TPv3 tunnel [[RFC4817](#)] or MPLS-in-UDP tunnel [[I-D.ietf-mpls-in-udp](#)]) could be used instead (For more details about this special usage, please refer to [[I-D.xu-spring-islands-connection-over-ip](#)]). Since the transport (i.e., the underlay) could be IPv4, IPv6 or even MPLS networks, the above approach of encoding the SFP information by an MPLS label stack is fully transport-independent which is one of the major requirements for the SFC encapsulation [[I-D.ietf-sfc-architecture](#)].

In addition, the service classifier could further impose metadata on the MPLS packet through the Network Service Header (NSH) [[I-D.quinn-sfc-nsh](#)] (As for how to contain the NSH within a MPLS packet, please see [Section 3.3](#)). Here the Service Path field within 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 a value indicating the total number of SFs within the service function path. The service index SHOULD be decreased by one on each SF or SF-proxy on behalf of the corresponding legacy SF. When the service index become zero, the NSH MUST be removed from the packet by the SF or SF-proxy on behalf of the corresponding legacy SF.

3.2. How to Contain Metadata within an MPLS Packet

Since the MPLS encapsulation has no explicit protocol identifier field to indicate the protocol type of the MPLS payload, how to indicate the presence of metadata (i.e., the NSH which is only used as a metadata container) in MPLS packets is a potential issue. There is a possible ways to address the above issue: SFFs allocate two different labels for a given SF, one indicates the presence of NSH while the other indicates the absence of NSH. This approach has no change to the current MPLS architecture but it would require more than one label binding for a given SF.

4. Acknowledgements

The authors would like to thank Loa Andersson and Andrew G. Malis for their valuable comments and suggestions on the draft. The authors would like to thank Adrian Farrel, Stewart Bryant, Alexander Vainshtein, Joel M. Halpern for their comments on how to indicate the presence of metadata within an MPLS packet.

5. IANA Considerations

TBD.

6. Security Considerations

TBD

7. References

7.1. Normative References

- [I-D.ietf-sfc-architecture]
Halpern, J. and C. Pignataro, "Service Function Chaining (SFC) Architecture", [draft-ietf-sfc-architecture-07](#) (work in progress), March 2015.
- [I-D.ietf-spring-segment-routing-mpls]
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Shakir, R., Tantsura, J., and E. Crabbe, "Segment Routing with MPLS data plane", [draft-ietf-spring-segment-routing-mpls-00](#) (work in progress), December 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.ietf-mpls-in-udp]
Xu, X., Sheth, N., Yong, L., Callon, R., and D. Black,
"Encapsulating MPLS in UDP", [draft-ietf-mpls-in-udp-11](#)
(work in progress), January 2015.
- [I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B.,
Litkowski, S., Horneffer, M., Shakir, R., Tantsura, J.,
and E. Crabbe, "Segment Routing Architecture", [draft-ietf-spring-segment-routing-01](#) (work in progress), February 2015.
- [I-D.quinn-sfc-nsh]
Quinn, P., Guichard, J., Surendra, S., Smith, M.,
Henderickx, W., Nadeau, T., Agarwal, P., Manur, R.,
Chauhan, A., Halpern, J., Majee, S., Elzur, U., Melman,
D., Garg, P., McConnell, B., Wright, C., and K. Kevin,
"Network Service Header", [draft-quinn-sfc-nsh-07](#) (work in
progress), February 2015.
- [I-D.xu-spring-islands-connection-over-ip]
Xu, X., Raszuk, R., Chunduri, U., and L. Contreras,
"Connecting MPLS-SPRING Islands over IP Networks", [draft-xu-spring-islands-connection-over-ip-04](#) (work in
progress), March 2015.
- [RFC4023] Worster, T., Rekhter, Y., and E. Rosen, "Encapsulating
MPLS in IP or Generic Routing Encapsulation (GRE)", [RFC 4023](#), March 2005.
- [RFC4817] Townsley, M., Pignataro, C., Wainner, S., Seely, T., and
J. Young, "Encapsulation of MPLS over Layer 2 Tunneling
Protocol Version 3", [RFC 4817](#), March 2007.

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/>

