INTERNET-DRAFT
Intended Status: Proposed Standard
Expires: January 1, 2015

X.Wei
Huawei Technologies

June 30, 2014

Re-classification analysis in SFC
draft-wei-sfc-re-classification-00

Abstract

Service Function Chaining (SFC) provides the ability to classify and steer a flow via some network service(s). Some traffic flows require re-classification to a new service chain. This may be, for example, the result of further analysis of initial packets, or detection of multiple types of media. This document discusses re-classification scenarios in SFC, and several deployment models for the re-classifier and relevant analysis are provided. The proposal will recommend some architectural constraints for the SFC design.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/1id-abstracts.html

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

Copyright and License Notice

Table of Contents

1 Introduction

Service Function Chaining (SFC) provides flexible configuration of services that are connected through the network. The requirements and problem space have been widely discussed in [PS], and several different SFC frameworks have been proposed in [Quinn],[Boucadair] and [Jiang] etc. These frameworks can be divided into four logical components: a centralized SFC controller, classifier, Service Function (SF) instance and SFC forwarding entity. The SFC controller is in charge of constructing service function chain for network traffic and certain network configurations; classifier is used to perform traffic classification, it classifies packets and adds a SFC ID, which is used to steer the packet along certain service chain, for each packet; SF instances are deployed to provide network service functions to traffics; SFC forwarding entities are in charge of forwarding packets to specific SF instances according to SFC ID contained in packets.

The concept of SFC ID is used to steering traffic along different service function chain, and additional information named as 'metadata' could also be used to convey information between SF instances or between classifier and SF instance.

In SFC network, classifier maintains < match rule, service chain > pairs for classifying traffic to different service chain, and two of the most important role of classifier is classifying packets based on matching rules and tagging the packets with appropriate service chain ID according to the classification result.

The criteria to classify traffic could be based on different kinds of information, including simple information such as 5-tuple in IP header field, or complex information such as the processing result of DPI function. Because the classifier functional entity is located at the entrance of SFC domain and all of the traffic enters SFC domain through classifier, and it's possible for classifier to deal with a large amount of traffic, so in order to reduce the load of classifier

and to accelerate the processing speed, the match ruling used by classifier should be simple, for example using 5-tuple of packet header as match rule. It would be better for classifier only to implement network logic in dealing with traffic; and leave other specific SFs to implement service logic.

Normally, classifier is deployed at the boundary of SFC domain, and traffic is classified when it enters the SFC domain. But in some cases, it is not possible to classify the traffic properly the first time traffic enters the SFC domain, because the classifier might not get sufficient information about the traffic and could only classify traffic coarsely based on some simple information. In other cases,

the traffic flow might need to be steered to a different service chain, for example, due to the processing result of a certain SF, after it first classified at the entrance of SFC domain. In these cases, a re-classification of the traffic flow is needed. By re-classification, it is meant that the service chain ID of the traffic is changed to a new one, after it enters SFC domain.

We name the classifier that implements re-classification action as re-classifier here.

In this document we discuss re-classification scenarios and related re-classifier deployment models in SFC domain.

1.1 Terminology

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].

Match rule: The rule that used by classifier to match the traffic to a specific service chain. For example, match rule could be in the form of 5-tuple in IP header.
SFE: Service Forwarding Entity. Forwarding entity in SFC domain, which forwards traffic along the service chain based on service chain ID.
SFC domain: A network that implements SFC.

2. Scenarios

This section provides several re-classification scenarios in SFC domain. Based on the scenarios discussed here, analysis of impacts of re-classification on SFC framework will be provided in the following section.

2.1 Case 1: Classifier lacks information about the traffic

When a flow enters SFC domain for the first time, SFC domain may have little or coarse-grained information, about the flow, e.g. which subscriber the flow belongs to, and more detailed, fine-grained information such as application type, security status or others may not be known. So classifier first classifies the packet according to coarse-grained information to a service chain, and the processing result of certain SF along the service chain could provide fine-grained information about the traffic, and then the re-classifier steers the traffic to another service chain based on the fine-grained information.

There are different SFs that could provide fine-grained information about the traffic, for instance, DPI (Deep Packet Inspection) can analyze the content of packet and determine the application type of the traffic, and then different kind of application traffic for the same subscriber could be steered to different service chain.

2.2 Case 2: Traffic bypasses certain SF

The use case of long-lived flow is introduced in [Krishnan]. B. After the long-lived flow has been processed by certain expensive or highly specialized SF, subsequent packets could be steered to bypass this SF in order to save processing resource. [Krishnan] also provides several examples of this kind of SF such as transparent FW, CDN (Content Delivery Network).

When traffic has to bypass a SF of the service chain, it means the traffic should go through a different service chain. It thus means that re-classification of the traffic is needed. Figure 1 shows an example for this scenario.

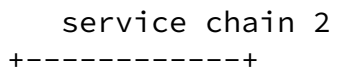

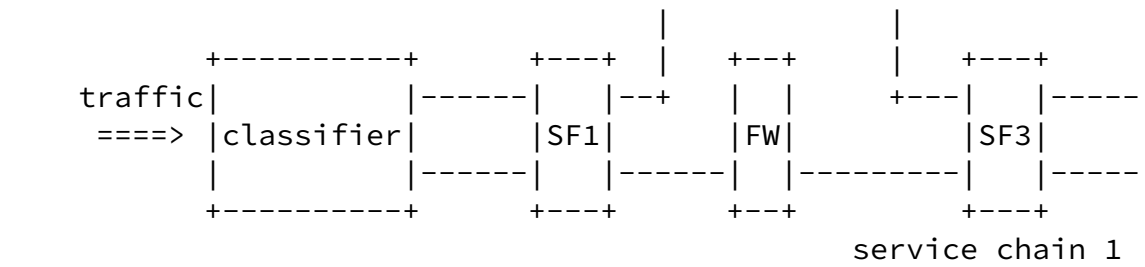

Figure 1 Traffic bypasses certain SF

In Figure 1, after a long-lived traffic flow enters the SFC domain for the first time, it is classified to service chain 1, which consists of a Firewall (FW). When the FW inspects the traffic and determines that it is from a trusted source, the traffic can subsequently bypass the FW. This requires re-classification that allows the flow to go through another service chain, e.g. service chain 2.

2.3 Case 3: Multiple content types in the same flow

In some applications, such as Web services using HTTP, different types of content may be transmitted in the same flow, i.e. belonging to the same TCP session. For example, when a user starts a TCP connection to YouTube and gets content using HTTP, there will be different kinds of media such as text, audio and video transmitted over the same TCP session.

For such flows, the content of different kinds might benefit from different network services, i.e. go through different service chain, for instance, in the same HTTP session, video may go through video optimizer, while text doesn't need any optimization. Though the classifier may determine the application type, e.g., based on port number, it not viable to assume that the classifier can distinguish different contents in traffic.

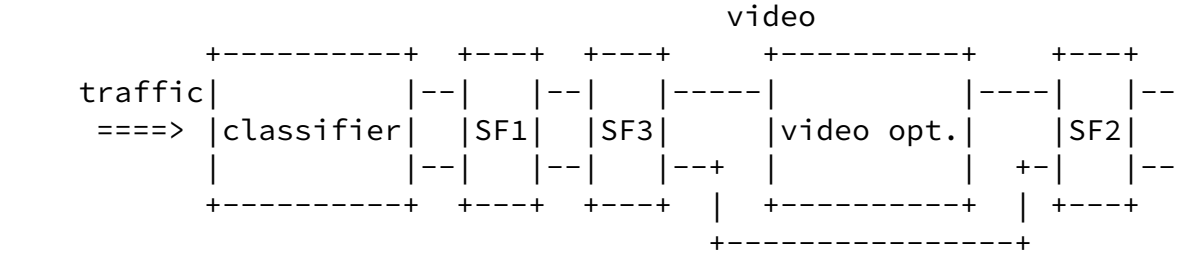

Figure 2 Example of re-classification of HTTP traffic

3 Deployment models of re-classifier

There are two different deployment models for the re-classification function, and this section will discuss these two models in detail.

3.1 Classifier plays the role of re-classifier

When traffic enters SFC domain through classifier, it will be classified to certain service chain based on matching rule configured in classifier. And then traffic is steered to specific SF instances along the service chain. In this sub-section, we discuss the re-classification mechanism that classifier plays the role of re-classifier and used to re-classify traffic to a different service chain, according to the procession result.

An overview description of this mechanism is depicted in Figure 3. If the processing result of certain SF would affect the service chain of a traffic flow, the SF reports processing result to controller to provide more detailed information about the traffic; after receiving traffic information from SF, controller could choose to form a new service chain for the traffic, and then controller updates the <match rule, service chain> pair in classifier; after updating of <match rule, service chain> pair, classifier uses the new <match rule, service chain> pair to classify the traffic. For example, in Figure 3, when traffic passes through classifier at the first time, it is classified to service chain 1 consisting SF1, SF2 and SF3 by classifier, the processing result of the traffic by SF2 would affect the classification behavior of classifier, and then classifier classifies the traffic to a different service chain 2 consisting of

SF4 and SF5.

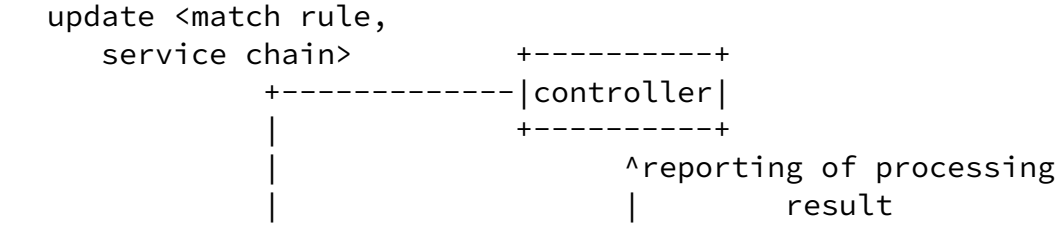

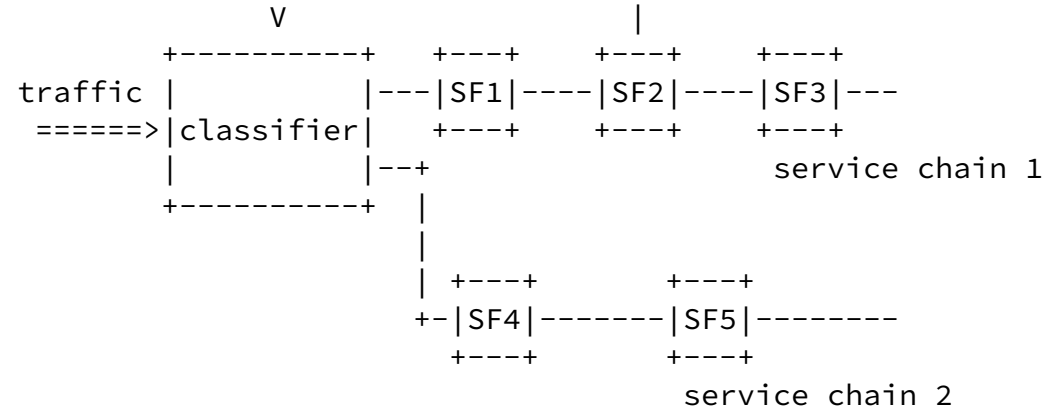

Figure 3: Classifier plays the role of re-classifier

Analysis:

- In this case, SF2 could be any kind of Service Function, such as Firewall, DPI etc, whose processing result of traffic flow could change the service chain of the traffic.
- This mechanism needs an interface to convey SF's processing result from SF to controller. The information conveyed from the SF to controller should at least include flow information and the processing result information. The protocol used in this interface could based on the existing protocols, such NetConf [RFC6241], Diameter [RFC6733] etc, or a new protocol might be needed.
- After the traffic flow is re-classified, the SF that cause the re-classification, i.e. SF2 in Figure 3, might not be included any more in the new service chain, so if the SF maintains state for the traffic flows state information, there should be a mechanism to release the state.

3.2 Re-classifier deployed independent from classifier

Besides the scheme described in sub-section 3.1, this sub-section provides another re-classification scheme in which re-classifier functionality is deployed independent from classifier which is deployed at the boundary of the SFC domain.

3.2.1 Re-classifier integrated with SF

In this case, the re-classifier is integrated with specific SF, e.g. FW, DPI etc, and it could re-classify traffic to a different

service chain according to the processing result of traffic by the

SF.

The traffic is first classified by classifier at the boundary of SFC domain, and then forwarded to specific SFs according to service chain ID encapsulated in traffic. When traffic goes through a SF which is integrated with re-classifier functionality, and the traffic needs to go through a different service chain, re-classifier could re-tag a new service chain ID in the traffic.

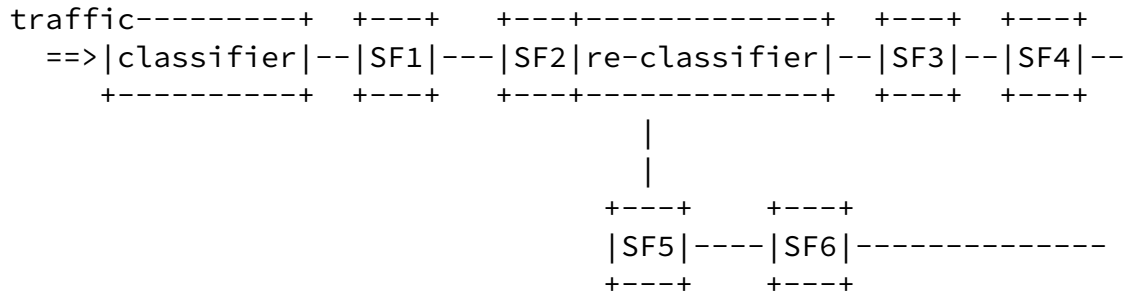

Figure 4 Re-classifier integrated with SF

Analysis:

- In this case, an interface between SFC controller and the SF that integrated with re-classifier functionality is needed to configure re-classifier with <match rule, service chain> pair. The protocol used in this interface could be the same as the interface between SFC controller and classifier.
- The re-classification in this deployment model only impacts the path that behind the SF which implements re-classification functionality.
- This scheme provides a more flexible choice to implement traffic re-classification, one or more re-classifier could be deployed in a SFC domain.

3.2.2 Re-classifier independent from SF

In cases that SF, which could impact service chain that the traffic goes through, is not integrated with a re-classifier, especially for legacy SF, a re-classifier independent from SF could be deployed.

For the re-classifier independent from SF there are also several deployment choices, for example, re-classifier could be deployed as an independent entity and attach to the output interface of specific SF; or re-classifier could be integrated with SFE.

When the traffic is processed by SF, the re-classifier could re-classify the traffic according to the output interface of SF or according to processing result information taken in the traffic, e.g.

the information in metadata.

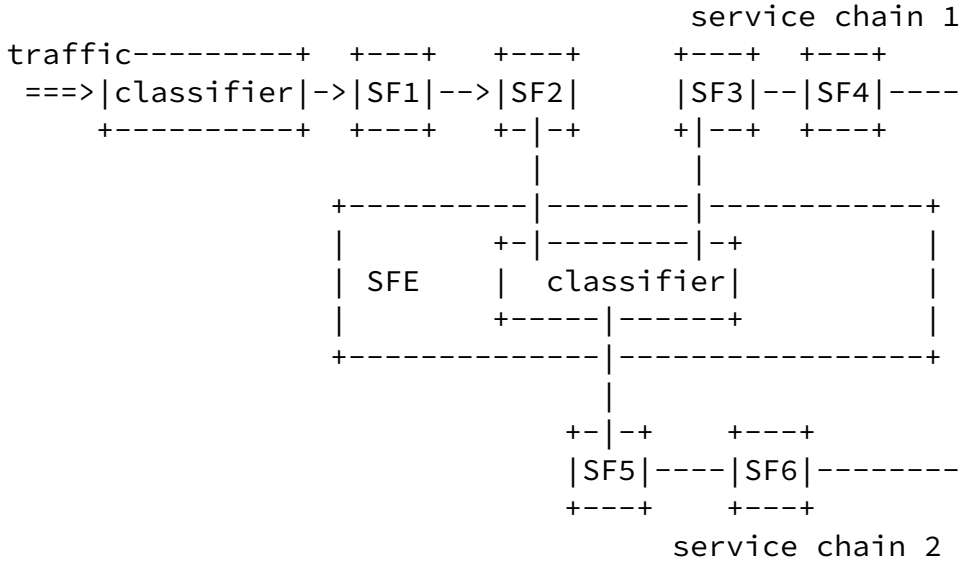

Figure 5: Re-classifier integrated with SFE

Analysis:

- In this case, SF is not required to support re-classification functionality.
- When integrated with SFE, the re-classifier could be shared by more than one SF.
- The processing result of SF should be conveyed in certain method to re-classifier, as discuss above the method used here could be judging from the output interface of SF or according to processing result information taken in the traffic, e.g. the information in metadata (in-band signal).
- An interface between SFC controller and re-classifier is needed to provision the re-classifier, if the re-classifier is integrated with SFE, then the interface between SFC controller and SFE could be used to exchange information between SFC controller and re-classifier; otherwise an independent interface between SFC controller and re-classifier should be exist.

4 IANA Considerations

This document requires no requirement for IANA.

5 Security Considerations

Security related issues is not involved.

6 Acknowledgments

Many thanks to John Kaippallimalil and Chunshan Xiong (Sam) for their valuable comments.

7 References

7.1 Normative References

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

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, June 2011.

[RFC6733] Fajardo, V., Ed., Arkko, J., Loughney, J., and G. Zorn, Ed., "Diameter Base Protocol", RFC 6733, October 2012.

7.2 Informative References

[Krishnan] R. Krishnan et al. "draft-krishnan-sfc-long-lived-flow-use-cases", April 21, 2014

[Quinn] P. Quinn et al. "draft-quinn-sfc-arch-05", May 5, 2014

[Boucadair] M. Boucadair et al. "draft-boucadair-sfc-framework-02", February 12, 2014

[Jiang] Y. Jiang. "draft-jiang-sfc-arch-01", February 14, 2014

Authors' Addresses

Xinpeng Wei
EMail: weixinpeng@huawei.com