

SFC WG
Internet-Draft
Intended status: Standards Track
Expires: May 3, 2017

W. Jing
C. Wang
ZTE Corporation
C. FONTAINE
Qosmos
October 30, 2016

Flow Release Notify
draft-jing-sfc-flow-release-notify-01

Abstract

Service function chain is the definition of an ordered set of service functions. After instantiated, the service function path is created and the classified traffic is steered through the corresponding service function path and then forwarded to the final destination. SFs and SFC Proxies do not know the termination of a service flow. This document describes a method to notify the SFs and SFC Proxies the termination of a service flow.

When one service flow goes through the SFP, there may create some states in some SFs or SFC Proxies. However, when the service flow terminates, the SFs and SFC Proxies are unaware of that and maintain the states as well. This document describes a method to notify the SFs and SFC Proxies the termination of a service flow and release the resources occupied by the flow.

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 May 3, 2017.

Copyright Notice

Copyright (c) 2016 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 |
| 2. | Convention and Terminology | 3 |
| 3. | Defination of Flow Termination Message | 3 |
| 3.1. | Data Plane Approach | 3 |
| 3.2. | Control Plane Approach | 4 |
| 4. | Controller Behavior | 5 |
| 5. | Classifier Behavior | 5 |
| 6. | Service Function Forwarder behavior | 6 |
| 7. | SFC Proxy Behavior | 6 |
| 8. | SFC-aware Service Function Behavior | 6 |
| 9. | Security Considerations | 6 |
| 10. | Acknowledgement | 7 |
| 11. | IANA Considerations | 7 |
| 12. | References | 7 |
| 12.1. | Normative References | 7 |
| 12.2. | Informative References | 7 |
| | Authors' Addresses | 7 |

[1.](#) Introduction

Service function chain is the definition of an ordered set of service functions. After instantiated, the service function path is created and the classified traffic is steered through the corresponding service function path and then forwarded to the final destination.

SFs and SFC Proxies allocate resources (e.g. memory) for service flow in order to process the packets of service flow correctly.

Typically, in current SFC deployment, there are many SFC-unaware SFs which need SFC Proxies to assist them to fulfill SFC forwarding.

When service flow goes through these SFC Proxies, there are states created which cost resources to assist the return traffic from the

SFC-unaware SFs to retrieve the original NSH encapsulation. When a service flow terminates, the corresponding states/resources should be cleaned up. Unfortunately, SFs and SFC Proxies do not know the termination of a service flow. As a result of that, they cannot release the resources immediately. Maybe one solution is to set lifetime for these states, but how long the lifetime should be set is an issue as well as what if the lifetime is asynchronous between different SFs and SFC Proxies.

This document tries to disclose this issue and describes a synchronous method to notify the SFC-aware SFs and SFC Proxies the termination of a service flow and then release the resource occupied by the service flow synchronously.

2. Convention and 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 [[RFC2119](#)].

The terms are all defined in [rfc7665](#) and [[I-D.ietf-sfc-nsh](#)].

3. Definition of Flow Termination Message

A message with Flow Termination Indicator is treated as flow termination message.

And also, the termination flow's identifier is also included in flow termination message. As what is the flow's identifier and how to define flow's identifier depends on specific scenarios. for example, flow's identifier could be 5-tuple in the packet, or the Flow ID, or the session ID, or something else which exclusively identify the flow.

When SFC components receive flow termination message, they MUST abstract the flow's identifier field in the receiving message and release the corresponding resource occupied by the flow.

As how SFC components receive the flow termination message, there have different approaches, including data plane approaches and control plane approaches. Here elaborates these approaches as follow.

3.1. Data Plane Approach

Here flow termination message is a NSH encapsulated message. It is generated by Classifier, and transported along the SFP and ended at

the end of SFP. An example of flow termination message is as figure 1.




Figure 1: Example of Flow Termination Message

The flow's identifier field in the NSH Payload of the flow termination message should have sufficient information to uniquely identify the flow that is terminated. Specifically, IP five tuples is a typical flow's identifier.

The FTI means Flow Termination Indicator. It is encapsulted in NSH.

There are several solutions to carry Flow Termination Indicator to indicate this message is an flow termination message, such as:

- 1)use a reserved bit in the Base header;
- 2)use a bit in the mandatory context header when MD type=1;
- 3)specify a Variable Context Headers when MD type=2.

3.2. Control Plane Approach

One of the control plane approaches here is that SFC controller should assist in triggering flow termination message through interfaces defined in [[I-D.ietf-sfc-control-plane](#)] to notify SFC components to release the resource.

Another control plane approach here is that SFC Classifier triggers a new SFC OAM message as flow termination message along the SFP.


```

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Ver|O|C|R|R|R|R|R|   Length   | MD-type=0x1 | OAM Protocol |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Service Path ID           | Service Index |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Mandatory Context Header           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Mandatory Context Header           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Mandatory Context Header           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Mandatory Context Header           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|FTI Msg Type |   Length   | Flow's Identifier(available |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 2: SFC OAM message as flow termination message

FTI Msg Type: TBD.

Len: depends on the length of Flow's Identifier.

Flow's Identifier: idenitifer the flow uniquely.

4. Controller Behavior

Under SFC controller scenario, if controller acquires the termination of flow, it may generate flow termination message and send it to the SFC components related to the flow. How the controller detects the termination of flow is out of the scope of this document.

5. Classifier Behavior

Under other scenarios, if classifier acquires the termination of flow, it may generate flow termination message and send it to the SFP of the flow. How the classifier detects the termination of flow is out of the scope of this document.

Here are some examples of how to detect the termination of flow:

1)in case of mobile network, classifier can be collocated with PGW. As per 3GPP specification, PGW has interfaces like S5/S8, Gx,Gy,S6b. All interfaces listed above can be used to detect the termination of flow. Specifically, Gx interface can be used by PGW to get online/offline information from PCRF.

2)Classifier may have DPI function, which can observe the TCP FIN packet which is termination signal of TCP flow. And so on.

6. Service Function Forwarder behavior

SFF treats flow termination message as normal traffic in service chain and forwards it according to SPI/SI.

But, under some circumstances, there may be some states maintained in the SFFs related to the flow. then, these states need to be released as well if such kind of SFFs receive flow termination message.

7. SFC Proxy Behavior

The proxy accepts packets from the SFF on behalf of the SF. It removes the SFC encapsulation, and then uses a local attachment circuit to deliver packets to SFC-unaware SFs. It also receives packets back from the SF, reapplies the SFC encapsulation, and returns them to the SFF for processing along the service function path. refer to [[RFC 7665](#)]

Thus, it is necessary for SFC proxy to maintain a state for each flow. When traffic is returned from the SFC-unaware SFs, SFC proxy reapplies the SFC encapsulation according to the encapsulation information stored in the states table. Such states consume a lot of memory, because millions of states would be maintained.

When SFC Proxy receives flow termination message, it should take action to release the resources of the flow, which is identified by the flow's identifier abstracted in the flow termination message. And then decrements service index and returns the flow termination message back to SFF.

8. SFC-aware Service Function Behavior

When SFC-aware SF receives flow termination message, it should take action to release the resources occupied the flow, which is identified by the flow's identifier. And then decrements service index and returns the flow termination message back to SFF.

9. Security Considerations

TBD

10. Acknowledgement

The editors would like to thank David Dolson for a thorough review and useful comments.

11. IANA Considerations

TBD

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", [RFC 7665](#), DOI 10.17487/RFC7665, October 2015, <<http://www.rfc-editor.org/info/rfc7665>>.

12.2. Informative References

- [I-D.ietf-sfc-control-plane] Boucadair, M., "Service Function Chaining (SFC) Control Plane Components & Requirements", [draft-ietf-sfc-control-plane-08](#) (work in progress), October 2016.
- [I-D.ietf-sfc-nsh] Quinn, P. and U. Elzur, "Network Service Header", [draft-ietf-sfc-nsh-10](#) (work in progress), September 2016.

Authors' Addresses

WeiDong Jing
ZTE Corporation
No.6 Huashen Rd, Yuhuatai District
Nanjing
China

Email: jing.weidong1@zte.com.cn

Cui Wang
ZTE Corporation
No.50 Software Avenue, Yuhuatai District
Nanjing
China

Email: wang.cui1@zte.com.cn

Christophe FONTAINE
Qosmos
8 rue Bernard-Bufferet
75017 Paris
France

Email: christophe.fontaine@qosmos.com

