SFC WG
Internet-Draft
Intended status: Standards Track
Expires: January 7, 2016

Ting. Ao
ZTE Corporation

July 6, 2015

Analysis of the SFC scalability
draft-ao-sfc-scalability-analysis-00.txt

Abstract

SFC as a chain of a set of service function, should be scalable to meet all kinds of requirements. The scalability of SFC means the SFC could be elastic to accomodate one or more SFs join the SFC , or leave the SFC. The document present four cases of the scalability, and analysis the data plane and the control plane to implement the scalable SFC.

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 January 7, 2016.

Table of Contents

1. Introduction

Service Function Chain (SFC) is the chain with a series of ordered Service Functions(SF). The SFC maybe changed because of load balance , failure, or other management requirement. We call it SFC's scalability. The SFC being scalable means that the Service Functions can be added or removed from the path of this SFC. With this capability, SFC is more flexible and elastic to adapt all kinds of requirements.

In this document, we will present four use cases on SFC scale-out and scale-in, and analysis some requirements to support such capability.

2. Terminology

SFC(Service Function Chain): An ordered set of some abstract SFs.

SFC Scale-out: One or more SFs are added into the path of the SFC for the sake of load balance, protection or other new services requirement.

SFC Scale-in: One or more SFs are removed from the path of the SFC for the sake of the SFs are by-passed or the SFs are failed.

3. Use cases

This clause is going to list four use cases to describe the scalability of the SFC.

3.1. Join

One or more new SFs are required to join the SFC for the traffic that has been classified to get more Service Functions to process. This case may be the opposite scenario of the by-pass. At this time, the SFC is scaled out.

When a SF is needed to join the SFC, control plane need to notify the previous SFF that a new SF joins the SFC as next SF and its next hop should be this SF. In this case, SFF forward the traffic not only according to the SFPID but also according to the metadata in the SFC header.

3.2. Redundancy

One or more SFs are added into the SFC for Redundancy or Load balance. This case is different with the first case (3.1) in that the SF is same with one of the SF that is on the path of the SFC. The new SF is used to protect the current corresponding SF or to offload the current corresponding SF. This is also a SFC scale-out.

In this case, control plane need to notify the previous SFF that a new SF joins the SFC as a redundancy SF and its next hop should be a group. To make sure the correctly forwarding, it's required that there is a flow id field in the SFC header so that SFF can select a SF from group according to the flow id.

3.3. By-pass

This is a SFC scale-in case. This use case has been described in [draft-ietf-sfc-long-lived-flow-use-cases] and [draft-kumar-sfc-offloads]. In these two draft, a SF is offload because it is not necessary to steer the traffic to the SF to improve the performance.

3.4. Fault

When SF in one SFC is failed out or removed out because of the no need of load balance or protection, the SFC is scaled in also.

For this case, it's also required that the previous SFF should be notified that its next hop should be changed to the next SF of the SF.

From above SFC scale-out and scale-in cases, we can get some requirements about control protocol that it should send out a message about next hop modification to SFF to support SFC dynamic scale

4. Data Plane

For the load balance or protection switch case of the SFC scale capability, it is required that there is a entropy field in the SFC head so that SFF can forward the traffic to different load balance SF according to this entropy field. The entropy field can be named as Flow ID which should be in SFC header.

This requires Classifier not only classifies the traffic to different SFPID, but also classifies the traffic with different Flow ID.

5. Control Plane

Control plane for SFC would be centralized or distributed.

5.1. Centralized CP

Controller is required to:

a) Send a message to SFF that the joined SF connected to set the correct SFPID and its next hop.

b) Send register message to previous SFF with some information. Such information not only includes next hop locator, but also includes an indicator that if the next hop is a new joined SF or the next hop is a new SF that added into a group. If the indicator is a new joined SF, it means a new SF will join the SFC. If the indicator is a group SF, it means a new SF will be added into a group for load balance or protection.

c) Send de-register message to previous SFF with some information. Such information not only includes next hop locator, but also includes an indicator that if the next hop is the next SF because the current SF is by-passed, or the next hop is the SF that is removed from a group. If the indicator is the by-passed SF, it means the current SF is by-passed or is leaving from the SFC. If the indicator is group SF, it means the current SF will be removed into a protection group that is for load balance or protection.

5.2. Distributed CP

Distributed CP can be used in Plug-and-Play scenario. Distributed CP requires:

a) The SF that needs to join the SFC or by-pass from the SFC should notify the SFF it connects by a message.

b) The SFF should send a register message to the previous SFF with some information. Such information not only includes next hop locator, but also includes an indicator that if the next hop is a new joined SF or the next hop is a new SF that added into a group. If the indicator is a new joined SF, it means a new SF will join the SFC. If the indicator is a group SF, it means a new SF will be added into a group for load balance or protection.

c) The SFF send de-register message to previous SFF with some information. Such information not only includes next hop locator, but also includes an indicator that if the next hop is the next SF because the current SF is by-passed, or the next hop is the SF that is removed from a group. If the indicator is the by-passed SF, it means the current SF is by-passed or is leaving from the SFC. If the indicator is group SF, it means the current SF will be removed into a protection group that is for load balance or protection.

6. Security Considerations

For the scalability of the SFC, security is very important to be considered. Before allow the SF to join to the SFC, it is required to check the SF's security firstly.

7. IANA Considerations

N/A

8. Information References

[I-D.ietf-sfc-architecture]
Halpern, J. and C. Pignataro, "Service Function Chaining (SFC) Architecture", draft-ietf-sfc-architecture-09 (work in progress), June 2015.

[I-D.ietf-sfc-nsh]
Quinn, P. and U. Elzur, "Network Service Header", draft-ietf-sfc-nsh-00 (work in progress), March 2015.

[RFC7498] Quinn, P. and T. Nadeau, "Problem Statement for Service Function Chaining", RFC 7498, April 2015.

Author's Address

Ting Ao
ZTE Corporation
No.50, Ruanjian Road
Nanjing 210012
China

Phone: +86 25 88016576
Email: ao.ting@zte.com.cn