SFC Netmod
Internet-Draft
Intended status: Standards Track
Expires: April 2, 2016

R. Penno
P. Quinn
C. Pignataro
Cisco Systems
D. Zhou
Intel Corporation

September 30, 2015

Services Function Chaining Traceroute
draft-penno-sfc-trace-03

Abstract

This document defines a protocol that checks the liveness and report the service-hops of a service path. .

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].

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 April 2, 2016.

Copyright Notice

Table of Contents

1. Introduction

This document defines a protocol that allows a user to check liveness and get reports of the service-hops of a service path

2. Definitions and Acronyms

The reader should be familiar with the terms contained in [I-D.ietf-sfc-architecture], ,[I-D.ietf-sfc-architecture] and [I-D.quinn-vxlan-gpe]

3. SFC Trace

A trace packet uses the same NSH header as MD-type 1 with a few differences: OAM Bit and Next Protocol.

SFC Trace Request packet format

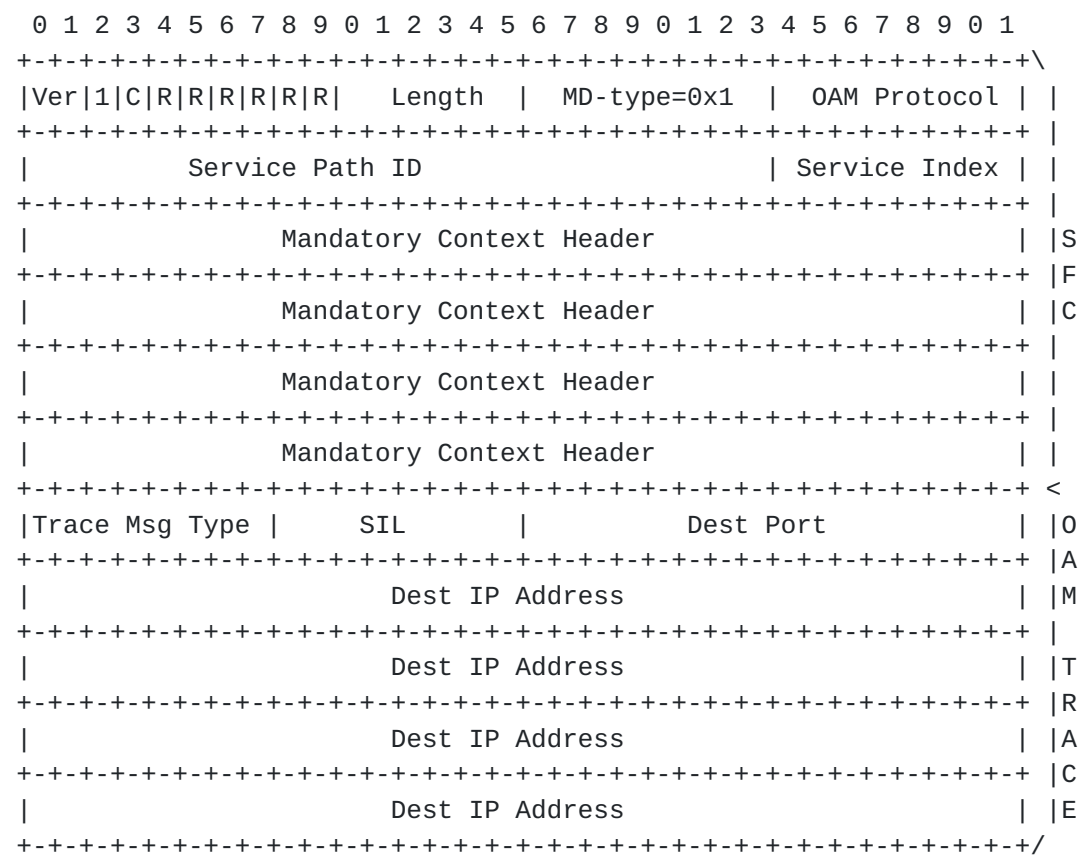

(postamble)

Ver: 1

OAM Bit: 1

Length: 6

MD-Type: 1

Next Protocol: OAM Protocol

Trace Msg Type: 1 for Trace Request and 2 for Trace Report

SIL: Service Index Limit: At least one less than the Starting Index

Dest Port: The trace report must be sent to this destination Port

Dest IP: the trace report must be sent to this destination IP address

For simplicity in building and parsing request and response packets, NSH Trace always uses fixed-size 128-bit IP address fields for both IPv6 addresses and IPv4 addresses.

When the address field holds an IPv6 address, the fixed-size 128-bit IP address field holds the IPv6 address stored as is.

When the address field holds an IPv4 address, an IPv4-mapped IPv6 address [RFC4291] is used (::ffff:0:0/96). This has the first 80 bits set to zero and the next 16 set to one, while its last 32 bits are filled with the IPv4 address. This is unambiguously distinguishable from a native IPv6 address, because an IPv4-mapped IPv6 address [RFC4291] would not be valid for a mapping.

When checking for an IPv4-mapped IPv6 address, all of the first 96 bits MUST be checked for the pattern -- it is not sufficient to check for ones in bits 81-96.

The all-zeros IPv6 address MUST be expressed by filling the fixed-size 128-bit IP address field with all zeros (::).

The all-zeros IPv4 address MUST be expressed by 80 bits of zeros, 16 bits of ones, and 32 bits of zeros (::ffff:0:0).

Allowing the client to insert the destination IP and port where it expects to receive reports in the NSH header allows for NAT traversal. In other words, if the client is behind a NAT, it can acquire a stable external IP:port and put as the destnation IP and port in the NSH header. This would allow NSH traceroute to function behind a NAT.

SFC Trace Report

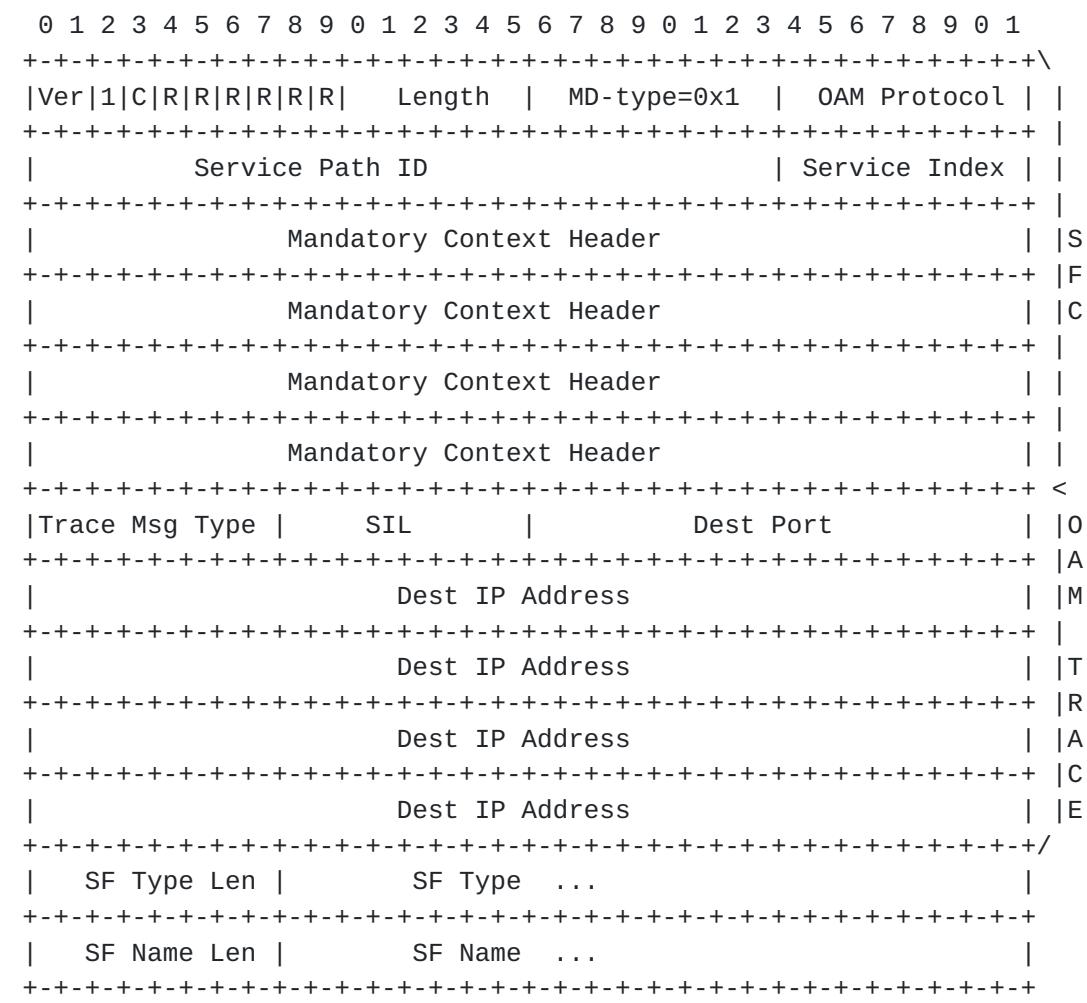

(postamble)

A trace report packet carries the identification of the Service Function that last processed the packet. In all other aspects it is exactly the same as a trace request.

SF Type Len: The Type Length in 4-byte words.

SF Type: A string representing the SF type padded to a 4-byte boundary and encoded with UTF-8. Service types can be found and registered in [I-D.penno-sfc-yang].

SF Name Len: The Name Length in 4-byte words.

SF Name: A string representing the Service Function padded to a 4-byte boundary and encoded with UTF-8. Service Function names and configuration can be found in [I-D.penno-sfc-yang].

4. Service Function Behavior

When a Service Function receives a SFC Trace request packet it performs the following actions:

1. Decrement Service Index

2. If Service Index is equal to the Services Index Limit add its identifying information at the end of the existing headers

3. Send packet back to Service Function Forwarder

5. Service Function Forwarder Behavior

A SFF will route trace packets based on service path ID and services index just like any other NSH packet. This guarantees that a trace packet follows the same path as data packets. The SFF will drop it and generate a report only in the following conditions:

- If the SI is equal or less than SIL

- If it can not find the next service-hop.

- If a SFF receives a trace packet with SI = 0.

In the cases enumerated above the SFF will proceed as following to build a trace report packet.

1. The SFF will use the same encapsulation as the received packet.

2. The destination IP:port will be the destination IP:port found in the OAM Trace NSH headers

3. The entire NSH +Trace Request headers + Report section will be copied from the received packet

4. The SFF will change the trace message type to trace report

If a SFF can not find the next service-hop for a trace packet, it will drop the packet and generate a report packet even if SIL is different from SI. This guarantees that the trace ends at the end of the path irrespective if SI has reached SIL or not. More importantly, it allow users to perform a trace that will traverse the entire path without having to know before hand the number of service-hops in the path by setting SIL to zero.

6. Implementation

SFC Trace was implemented in the Opendaylight projects and output of a 3 service-hop network can be found below.

```
sff_client.py --remote-sff-ip 10.0.1.41 --remote-sff-port 4789 --sfp-id 22 --
sfp-index 255 --trace-req --num-trace-hops 3

Sending Trace packet to Service Path and Service Index: (22, 255)
Trace response...
Service-hop: 0. Service Type: dpi, Service Name: SF1, Address of Reporting SFF:
('10.0.1.41', 4789)
Service-hop: 1. Service Type: firewall, Service Name: SF4, Address of Reporting
SFF: ('10.0.1.42', 4789)
Service-hop: 2. Service Type: napt44, Service Name: SF5, Address of Reporting
SFF: ('10.0.1.43', 4789)
Trace end
```

Implementation guideline for the client: If the trace request has a service index limit that would put the end of the trace beyond the service path, for example, starting Index=255, SIL=252 but only 2 service-hops in the path, the last trace response will have no report information. This is because no SF would detect that it is the end of the trace and include a report information

7. SFC Reverse Trace

Tracing a reverse path by sending a packet to the forward path is not always possible. The reason is that the sets of SFFs used in the forward and reverse might not have common elements.

8. IANA Considerations

OAM Protocol Type and a OAM protocol Message type.

9. Security Considerations

10. Acknowledgements

11. Changes

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>.

[RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, DOI 10.17487/RFC2616, June 1999, <http://www.rfc-editor.org/info/rfc2616>.

12.2. Informative References

[I-D.ietf-sfc-architecture]
Halpern, J. and C. Pignataro, "Service Function Chaining (SFC) Architecture", draft-ietf-sfc-architecture-11 (work in progress), July 2015.

[I-D.penno-sfc-yang]
Penno, R., Quinn, P., Zhou, D., and J. Li, "Yang Data Model for Service Function Chaining", draft-penno-sfc-yang-13 (work in progress), March 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.quinn-vxlan-gpe]
Quinn, P., Manur, R., Kreeger, L., Lewis, D., Maino, F., Smith, M., Agarwal, P., Yong, L., Xu, X., Elzur, U., Garg, P., and D. Melman, "Generic Protocol Extension for VXLAN", draft-quinn-vxlan-gpe-04 (work in progress), February 2015.

Authors' Addresses

Reinaldo Penno
Cisco Systems
170 West Tasman Dr
San Jose CA
USA

Email: repenno@cisco.com

Paul Quinn
Cisco Systems
170 West Tasman Dr
San Jose CA
USA

Email: paulq@cisco.com

Carlos Pignataro
Cisco Systems
170 West Tasman Dr
San Jose CA
USA

Email: cpignata@cisco.com

Danny Zhou
Intel Corporation
2200 Mission College Blvd.
Santa Clara CA
USA

Email: danny.zhou@intel.com