Network Work group
Internet-Draft
Intended status: Standards Track
Expires: September 14, 2017

N. Kumar
C. Pignataro
P. Quinn
Cisco Systems, Inc.
March 13, 2017

# IPFIX Information Element extension for SFC
## draft-kumar-ipfix-sfc-extension-05

Abstract

Service Function Chaining (SFC) is an architecture that enables any operator to apply selective set of services by steering the traffic through an ordered set of service functions without any topology dependency.

This document defines the required Information Elements to represent the details about service flows over any Service Function Path.

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 14, 2017.

Table of Contents



## 1. Introduction

[RFC7665] introduces and explains SFC architecture that enables any operator to apply selective set of services by steering the traffic through an ordered set of service functions without any topology dependency. Such ordered set of service functions to be applied to a packet is defined as service function chaining. As defined in [I-D.ietf-sfc-nsh], a classifier will add Network Service Header (NSH) to a packet that defines the corresponding service path to follow.

This document defines the required Information Elements to represent the details about traffic flows over any Service Function Path and export to Collector.

## 2. Requirements notation

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].

## 3. Terminology

This document uses the terminologies defined in [RFC7665] and [RFC7011]. In addition, this document defines the below terminologies:

Service Flow

A Service Flow is defined as a set of packets over a specific Service Function Path.

## 4. Network Service Header

Section 3.1 of [I-D.ietf-sfc-nsh] defines the Network Service Header format used by the classifier to encapsulate the traffic, carrying instruction about the service functions to be applied to the packet. This header comprises a 4 byte base header followed by a 4 byte service path header and a variable size context header.

In order to accomodate different needs from different use cases, there are 2 types of Network Service Header defined in [I-D.ietf-sfc-nsh] that preserves same Base header and Service Path header while differs in Context header. NSH MD-type 1 have a fixed size Mandatory Context header while NSH MD-type 2 have a variable size TLV based context header. The details are below:

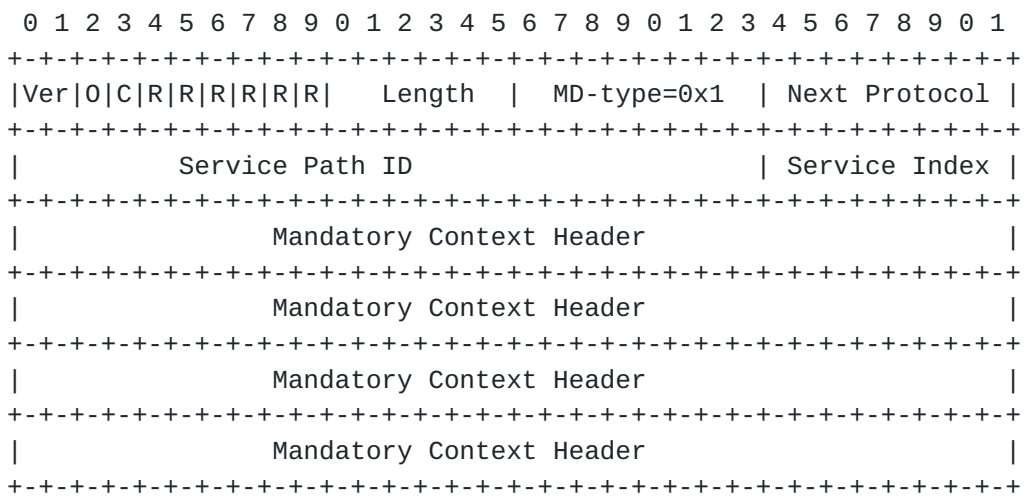


Figure 1: NSH MD-type 1

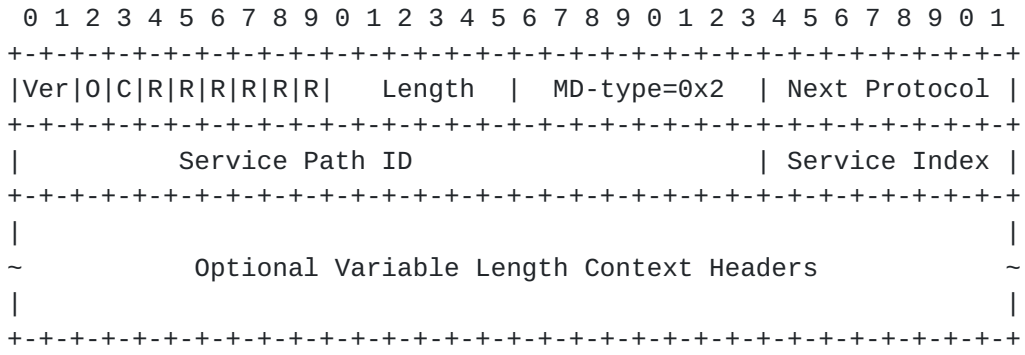


Figure 2: NSH MD-type 2

The details about different header fields are detailed in Section 3.4 and 3.5 of [I-D.ietf-sfc-nsh].

## 5. Flow measurement in SFC environment

SFC introduces the concept of steering user traffic over an ordered set of service function by utilizing service overlay between service functions over the existing network topology. The measurement of Service flow over Service Function Chain are required for various application such as but not limited to below:

- Capacity Planning - To ensure distributing load between Service Functions.

- OAM and troubleshooting - To measure the performance and troubleshooting failures.

- Traffic Profiling - Determine the characteristics of traffic flow over SFP.

- Security - To identify DoS attack or malicious intrusion.

In SFC environment, Classifier and Service Function Forwarder (SFF) are the different nodes that handles SFC encapsulation and it is appropriate to collect the Service Flow records in these nodes.

## 5.1. Observation Point

An Observation point in SFC environment is where Flow record for Service Flow will be collected and exported to the Collector. In a Classifier or SFF, an Observation point can be any physical or logical port that:

- Forwards NSH encapsulated packets or frame from Classifier to SFF.

- Receives NSH encapsulated packets or frame from Classifier or previous SFF.

- Receives NSH encapsulated packets or frame from Service Function after packet treatment applied.

- Forwards NSH encapsulated packets or frame to next Service Function Forwarder.

- Forwards NSH encapsulated packets or frame to Service Function for packet treatment.

## 5.2. Flow measurement

The ability to collect Flow record for different flows observed at the above range of Observation point allows an Operator to measure flow properties before and after the application of any service function within a service function path. An implementation SHOULD support the use of Information Elements defined in section 6 to measure and export the flow information. In addition, it also MAY support the use of other Flow keys relevant to the underlay network to collect any additional information from transport header encapsulating NSH header.

## 6. Service Flow Information Elements

This document defines the below set of Information Elements that are necessary for enabling IPFIX traffic measurement for Service Flow:

| ID | Name |
|---|---|
| TBD1 | nshBaseVersion |
| TBD2 | nshBaseFlags |
| TBD3 | nshBaseHeaderLength |
| TBD4 | nshBaseMDType |
| TBD5 | nshBaseNextProtocol |
| TBD6 | nshSphServicePathID |
| TBD7 | nshSphServiceIndex |
| TBD8 | nshMetadataMch |
| TBD9 | nshMetadataVch |
| TBD10 | nshIPv4NextSFF |
| TBD11 | nshIPv6NextSFF |
| TBD12 | nshEtherNextSFF |

### 6.1. nshBaseVersion

Description:

The Version field in NSH header.

Abstract Data Type: unsigned8

Element ID: TBD1

Data Type Semantic: identifier

Range: The valid range is 0-3.

Reference:

See Section 3.2 of [I-D.ietf-sfc-nsh]

### 6.2. nshBaseFlags

Description:

The flag bits from bit position 2 to 9 in NSH Base header. This information is encoded as a bit field.

Abstract Data Type: unsigned8

Element ID: TBD2

Data Type Semantic: flags

Reference:

See Section 3.2 of [I-D.ietf-sfc-nsh]

### 6.3. nshBaseHeaderLength

Description:

The length of the NSH header including any optional variable TLVs.

Abstract Data Type: unsigned8

Element ID: TBD3

Range: The valid range is 0-255.

Reference:

See Section 3.2 of [I-D.ietf-sfc-nsh]

### 6.4. nshBaseMDType

Description:

Defines the Metadata format beyond the NSH Base header.

Abstract Data Type: unsigned8

Element ID: TBD4

Data Type Semantic: identifier

Reference:

See Section 3.2 of [I-D.ietf-sfc-nsh]

### 6.5. nshBaseNextProtocol

Description:

This indicates the type of the payload packet encapsulated within the NSH header.

Abstract Data Type: unsigned8

Element ID: TBD5

Data Type Semantic: identifier

Reference:

See Section 3.2 of [I-D.ietf-sfc-nsh]

### 6.6. nshSphServicePathID

Description:

Service Path ID uniquely identifies the Service Chain which is a sequence of service function to be applied on the payload packet.

Abstract Data Type: unsigned32

Element ID: TBD6

Data Type Semantic: identifier

Reference:

See Section 3.3 of [I-D.ietf-sfc-nsh]

### 6.7. nshSphServiceIndex

Description:

Service Index identifies the next service function to be applied in the service chain.

Abstract Data Type: unsigned8

Element ID: TBD7

Range : The valid range is between 0-255.

Reference:

See Section 3.3 of [I-D.ietf-sfc-nsh]

### 6.8. nshMetadataMch

Description:

When MD Type is 1 on NSH header, Service Base header is followed by fixed size Mandatory Context Header. The format of this header varies depending on the implementation. This information element is of 16 bytes size.

Abstract Data Type: OctetArray

Element ID: TBD8

Reference:

See Section 3.4 of [I-D.ietf-sfc-nsh]

### 6.9. nshMetadataVch

Description:

When MD Type is 2 on NSH header, Service Base header is followed by Variable size Context Header. The format of this header varies depending on the implementation. This Informational element carries n octets from the NSH Service Path header.

A value of 64 reduced from nshBaseHeaderLength expresses how much Metadata was observed, while the remainder is padding.

Abstract Data Type: OctetArray

Element ID: TBD9

Reference:

See Section 3.5 of [I-D.ietf-sfc-nsh]

### 6.10. nshIPv4NextSFF

Description:

This defines the IPv4 address of the next SFF in the Service Function Path. This Information element is of size 4 bytes.

Abstract Data Type: ipv4Address

Element ID: TBD10

Data Type Semantic: identifier

Reference:

See Section 7.1 of [I-D.ietf-sfc-nsh]

### 6.11. nshIPv6NextSFF

Description:

This defines the IPv6 address of the next SFF in the Service Function Path. This Information element is of size 16 bytes.

Abstract Data Type: ipv6Address

Element ID: TBD11

Data Type Semantic: identifier

Reference:

See Section 7.1 of [I-D.ietf-sfc-nsh]

### 6.12. nshEtherNextSFF

Description:

This defines the Ethernet Address of the next SFF in the Service Function Path. This Information element is of size 16 bytes.

Abstract Data Type: macAddress

Element ID: TBD12

Data Type Semantic: identifier

Reference:

See Section 7.1 of [I-D.ietf-sfc-nsh]

## 7. IANA Considerations

To be Updated.

## 8. Security Considerations

TBD

## 9. Acknowledgement

The authors would like to thank Jim Guichard, Stewart Bryant, Benoit Claise, Richard Furr and Joel M. Harpern for their contribution.

## 10. Contributing Authors

TBD

## 11. References

### 11.1. Normative References

[I-D.ietf-sfc-nsh]
Quinn, P. and U. Elzur, "Network Service Header", draft-ietf-sfc-nsh-12 (work in progress), February 2017.

[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>.

[RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <http://www.rfc-editor.org/info/rfc7011>.

### 11.2. Informative References

[RFC7012] Claise, B., Ed. and B. Trammell, Ed., "Information Model for IP Flow Information Export (IPFIX)", RFC 7012, DOI 10.17487/RFC7012, September 2013, <http://www.rfc-editor.org/info/rfc7012>.

[RFC7013] Trammell, B. and B. Claise, "Guidelines for Authors and Reviewers of IP Flow Information Export (IPFIX) Information Elements", BCP 184, RFC 7013, DOI 10.17487/RFC7013, September 2013, <http://www.rfc-editor.org/info/rfc7013>.

[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>.

## Authors' Addresses

Nagendra Kumar
Cisco Systems, Inc.
7200 Kit Creek Road
Research Triangle Park, NC 27709
US

Email: naikumar@cisco.com

Carlos Pignataro
Cisco Systems, Inc.
7200 Kit Creek Road
Research Triangle Park, NC 27709-4987
US

Email: cpignata@cisco.com

Paul Quinn
Cisco Systems, Inc.
170 West Tasman Dr
San Jose, CA
US

Email: paulq@cisco.com