SFC Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 26, 2017

A. Farrel
Juniper Networks
L. Yong
Huawei USA
J. Drake
Juniper Networks

January 22, 2017

Operating the Network Service Header with Next Protocol "None"
draft-farrel-sfc-convent-00

Abstract

This document describes the use of the Network Service Header (NSH) in a Service Function Chaining (SFC) overlay network with no payload data and only carrying metadata. This is achieved by defining a new "next protocol" type value of "None".

This document illustrates some of the functions that may be achieved or enhanced by this mechanism, but it does not provide an exhaustive list of use cases, nor is it intended to be definitive about the functions it describes. It is expected that other documents will describe specific use cases in more detail and will define the protocol mechanics for each use case.

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 [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 July 26, 2017.

Copyright Notice

Copyright (c) 2017 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

The architecture for Service Function Chaining (SFC) is presented in [RFC7498]. That architecture enables packets to be forwarded along Service Function Paths (SFPs) to pass through various Service Functions (SFs) that act on the packets. This is achieved by inserting a Network Service Header (NSH) [I-D.ietf-sfc-nsh] into each packet to identify the SFP that the packet travels along (by means of a Service Path Identifier - SPI) and the point along the SFP

that the packet has reached (by means of a Service Index - SI). The SPI and SI are fields encoded in the NSH.

Packets are classified on entry to the SFC overlay network and have an NSH applied to them. Such packets are forwarded between Service Function Forwarders (SFFs) and each SFF may hand the packet off to one or more SFs according to the definition of the SFP.

The packet classifier or SFs may wish to share information (possibly state information) about the SFP, the traffic flow, or a spcific packet, and they may do this by adding "metadata" to packets as part of the NSH. Metadata may be used to enhance or enable the function preformed by SFs, may enable coordination between SFs, or may be used to assist a network operator in the diagnosis and monitoring of an SFP.

This document defines a mechanism for metadata to be carried on an SFP without the need for payload data. This may enable diagnosis and monitoring of SFPs, and coordination between SFs, without the need for traffic to be flowing, and without the need to rewrite data packets to insert what might be substantial amounts of metadata.

This function is achieved by defining a new vlue for the NSH "Next Protocol" field to indicate "None".

This document illustrates some of the functions that may be achieved or enhanced by this mechanism, but it does not provide an exhaustive list of use cases, nor is it intended to be definitive about the functions it describes. It is expected that other documents will describe specific use cases in more detail and will define the protocol mechanics for each use case.

2. The Network Service Header

The NSH is defined in [I-D.ietf-sfc-nsh]. It includes a field called "Next Protocol" that is used to indicate the nature of the payload data that follows the NSH. The field can be used by any component that processes the NSH (for example, to understand how to interpret and parse the payload) and by nodes at the end of the SFP that remove the NSH and forwrd the payload data.

2.1. Next Protocol None

This document defines a new value for the Next Protocol field. When set to TBD1, the field indicates that the next protocol is "None" meaning that there is no user/payload data following the NSH.

3. Processing Rules

A node wishing to send metadata without a data packet MUST:

- Create a packet carrying an NSH and the desired metadata
- Set the Next Protocol field to TBD1
- Ensure that there are no bytes following the end of the NSH
- Encapsulate and send the packet as normal for the SFP.

Note that a packet with no payload data may be simply inserted at the head end of an SFP (such as by a classifier) and may be easily forwarded by an SFF or SF on the SFP using the normal processing rules defined in [I-D.ietf-sfc-nsh]. Such a packet may only be inserted into the middle of an SFP by a node that knows (by control plane or management plane means) the correct values of SPI and SI to use at that point on the SFP.

A transit node (SFF, SF, or classifier) receiving a packet with Next Protocol indicating "None" MUST NOT attempt to parse or process beyond the end of the NSH, but can process the NSH and especially the metadata as normal.

A node that is the egress of an SFP would normally strip the NSH and forward the payload according to the setting of the Next Protocol field. Such nodes MUST NOT attempt to forward the paylpoad of packets with Next Protocol indicating "None". These packets would, in any case be zero length.

4. Backward Compatiblity

Nodes that do not understand the meaning of Next Protocol set to "None" will be unable to parse the payload data just as they would be unable to process the payload if the Next Protocol field was set to any other value they do not understand. Such nodes MAY, according to normal behavior of [I-D.ietf-sfc-nsh], perform protocol independent processing of the payload (such as hashing the whole packet) constrained by knoweledge of the packet length. If the packet length is not known, such processing obviously cannot be performed.

Nodes at the end of an SFP possibly forward packets with no knowledge of the payload in a "pop and forward" form of processing where the NSH is removed and the packet is simply put on an interface. It is a general processing rule for all forwarders that they SHOULD NOT attempt to send packets with zero length, and this will be the case when the NSH Next Protocol is "None".

5. Overview of Use Cases

5.1. Per SFP Metadata

Per SFP metadata may be sent along the path of an SFP simply by setting the correct SPI in the NSH, and setting the SI to the correct value for the introduction of the metadata. Classifiers will know the correct SI value for the point in the SFP at which they reside.

5.2. Per Flow Metadata

Per flow metadata is complicated if there is more than one flow carried on an SFP. If there is just one flow on an SFP then there is no difference between per-flow metadata and per-SFP metadata.

In normal processing, the flow to which per-flow metadata applies can be deduced by looking at the payload data in the context of the value of the Next Protocol field. When Next Protocol indicates "None" this cannot be done. In this case the identity of the flow would need to be carried in the metadata.

5.3. Coordination Between SFs

A pair of SFs (adjacent or not) on an SFP may desire to coordinate state and may do this by sending information encoded in metadata.

To do this using the mechanisms defined in this document:

- There must be an SFP that passes through the two SFPs in the direction of sender to receiver
- The sender must know the correct SPI to use
- The sender must know the correct SI to use for the point at which it resides on the SFP
- Ideally the receiver will know to remove the packet from the SFP and not forward it further as this might share metadata wider than desirable and would cause unneccessary packets in the network. Note, however, that continued forwarding of such packets would not be substantially harmful in its own right.

Note that technically (according to the SFC architecture) the process of inserting a packet into an SFP is performed by a Classifier. However, a Classifier may be co-resident with an SF.

Note also that a system with SFs that need to coordinate between each other may be configured so that there is a specific, dedicated SFP between those service functions that is used solely for this purpose.

5.4. Operations, Administration, and Maintenance (OAM)

Requirements for Operations, Administration, and Maintenance (OAM) in SFC networks are discussed in [I-D.ietf-sfc-oam-framework]. The NSH definition in [I-D.ietf-sfc-nsh] includes an O-bit that indicates that packet contains OAM information.

Since OAM information will be carried in packets that also include payload data, that information must be carried in metadata. Therefore, the mechanism defined in this document can be used to carry OAM information independent of payload data.

Sending OAM separate from (but interleaved with) packets that carry payload data may have several advantages including:

- Sending OAM when there is no other traffic flowing.
- Sending OAM at predictble intervals.
- Measuring paths qualities distinct from bevaior of SFs.
- Sending OAM without needing to rewrite payload data buffers.
- Keeping OAM processing components separate from other processing components.

5.5. Control Plane and Management Plane Uses

As described in Section 5.3 SFPs can be established specically to carry metadata-only packets. And as described in Section 5.1, metadata-only packets can be sent down existing SFPs. This means that metadata-only packets can be used to carry control plane and management plane messages used to control and manage the SFC network.

In effect, SFPs can be established to serve as a Data Control Network (DCN) or Management Control Network (MCN).

5.6. Non-Applicable Use Cases

The mechanisms described in this document are not applicable to per-packet metadata.

6. Security Considerations

Metadata-only packets as enabled by this document create a covert channel. However, this is only different from the same feature in the normal NSH in that it can be sent without the presence of a data flow.

Metadata may, of course, contain sensitive data and may also contain information used to control the behavior of SFs in the network. As such, this data needs to be protected according to its value and according to the perceived vulnerabilities of the network. protection of metadata may be achieved by using encrypted transport between SFC entities or by encrypting the metadata in its own right. The need to protect the metadata is not modified by this document.

The mechanism desrcibed in this document might possibly be used to introduce packets into the SFC overlay network. Therefore measures SHOULD be taken to ensure authorization of sources of such packets, and tunneling of such packets into the network SHOULD be prevented.

Further discussion of NSH security is presented in [I-D.ietf-sfc-nsh].

7. IANA Considerations

IANA has been requested to create a registry of "Next Protocol" values in [I-D.ietf-sfc-nsh]. This document requests IANA to allocate a value from that registry to indicate "None" (TBD1 in this document).

It is strongly suggested that a value of 0 (zero) be assigned.

8. Contributors

TBD

9. Acknowledgements

Thanks to the attendees at the SFC interim meeting in Westford in January 2017 for discussions that suggested the value of this document.

10. References

10.1. Normative References

[I-D.ietf-sfc-nsh]
Quinn, P. and U. Elzur, "Network Service Header", draft-ietf-sfc-nsh-10 (work in progress), September 2016.

[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>.

10.2. Informative References

[I-D.ietf-sfc-oam-framework]
Aldrin, S., Krishnan, R., Akiya, N., Pignataro, C., and A. Ghanwani, "Service Function Chaining Operation, Administration and Maintenance Framework", draft-ietf-sfc-oam-framework-01 (work in progress), February 2016.

[RFC7498] Quinn, P., Ed. and T. Nadeau, Ed., "Problem Statement for Service Function Chaining", RFC 7498, DOI 10.17487/RFC7498, April 2015, <http://www.rfc-editor.org/info/rfc7498>.

Authors' Addresses

Adrian Farrel
Juniper Networks

Email: afarrel@juniper.net

Lucy Yong
Huawei USA
5340 Legacy Dr.
Plano, TX 75024
US

Phone: +1 858 6511 4478
Email: lucy.yong@huawei.com

John Drake
Juniper Networks

Email: jdrake@juniper.net