SFC Working Group A. Farrel
Internet-Draft J. Drake
Intended status: Standards Track Juniper Networks
Expires: December 31, 2017 June 29, 2017

Operating the Network Service Header (NSH) with Next Protocol "None"
draft-farrel-sfc-convent-02

Abstract

This document describes the use of the Network Service Header (NSH) in a Service Function Chaining (SFC) enabled network with no payload data and carrying only metadata. This is achieved by defining a new NSH "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 December 31, 2017.

Table of Contents

1. Introduction

An architecture for Service Function Chaining (SFC) is presented in [RFC7665]. That architecture enables packets to be forwarded along Service Function Paths (SFPs) to pass through various Service Functions (SFs) that act on the packets. Each packet is encapsulated with a Network Service Header (NSH) [I-D.ietf-sfc-nsh] identify the SFP that the packet travels along (by means of a Service Path Identifier - SPI) and the hop (i.e., the next SF to be executed)

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 at the SFC ingress boundaries (section 4.4 of [RFC7665]) and have an NSH applied to them. Such packets are forwarded between Service Function Forwarders (SFFs) using tunnels across the underlay network, and each SFF may hand the packet off to one or more Service Function Instances (SFIs) according to the definition of the SFP.

The SFC classifier or any SFC-aware SFI may wish to share information (possibly state information) about the SFP, the traffic flow, or a specific 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 SFC-aware SFs, may enable coordination and data exchange between SFIs, or may be used to assist a network operator in the diagnosis and monitoring of an SFP. The nature of metadata to be supplied and consumed is implementation- and deployment-specific.

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 SFC-aware SFIs, 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 value for the NSH "Next Protocol" field to indicate "None". Such packets are contained within the SFC-enabled domain.

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 forward 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.

When the next protocol is "None" the rest of the NSH still has meaning and, in particular, the metadata carried in the NSH may still be present.

3. Processing Rules

An SFC-aware node wishing to send metadata without a data packet:

- MUST create a packet carrying an NSH and the desired metadata
- MUST set the "Next Protocol" field to TBD1
- SHOULD ensure that there are no bytes following the end of the NSH (i.e., that there is no payload data)
- MUST encapsulate and send the packet as normal for tunneling to the next hop on the SFP as normal for an NSH packet.

A packet with no payload data may be simply inserted at the head end of an SFP (such as a Classifier) and may be easily forwarded by an SFF or SFI on the SFP using the normal processing rules defined in [I-D.ietf-sfc-nsh].

A packet with no payload may also be generated by an SFC-aware SFI as a result of processing an incoming packet (i.e., triggered by a condition arising from processing a normal NSH packet with a payload). In such cases, the SPI/SI can be inherited from the original packet or can be set according to information supplied through the control plane or management plane. This document does not further specify the triggers to generate an NSH packet with a "Next Protocol" set to "None".

A transit node (SFF, SFI, 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 forward packets with "Next Protocol" indicating "None" even if there some bytes after the NSH.

4. Backward Compatibility

This section describes procedures for default handling on unknown "Next Protocol" field values. This material updates the procedures described in [I-D.ietf-sfc-nsh] and may be transferred to that document.

SFC-aware nodes that do not understand the meaning of a value contained in the "Next Protocol" field of the NSH are unable to parse the payload. Such nodes are not obliged to discard the packet unless they are specifically called upon to be able to examine the payload.

Thus:

- Transit SFFs will normally not inspect the "Next Protocol" field or the packet payload and will forward the packets based solely on the SPI/SI

- An SFC Proxy must not pass to an SFI a packet of type where it cannot indicate the packet type to the SFI

- An SFC Proxy must not pass to an SFI a packet of type that the SFI does not support

- An SFC Proxy should not return to the SFF a packet it has not passed to the SFI

- An SFI should not return to the SFF a packet it hasn’t processed unless local policy defines "process" for this SF to mean "do not process" in this case.

- Reclassifiers would normally require to understand the payload packet type, but it is possible to imagine reclassifiers that take action based on unknown values of the "Next Protocol" field or that perform protocol-independent actions (such as hashing the whole packet).

All this means that legacy SFC-aware nodes that are unaware of the meaning of the "Next Protocol" value "None" will act as follows:

- SFFs will forward the packets

- SFC Proxies will drop the packets

- SFIs will most likely drop the packets

- Reclassifiers will most likely drop the packets

SFC-aware 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 and the payload protocol is known a priori (or assumed). It is a general processing rule for all forwarders that they SHOULD NOT attempt to send packets with zero length, and since packets with the NSH "Next Protocol" set "None" are expected to have zero payload length.

5. Overview of Use Cases

5.1. Per-SFP Metadata

Per-SFP metadata is metadata that applies to an SFP and any data packets on that SFP. It does not need to be transmitted with every packet, but can be installed at the SFIs on the SFP and applied to all packets on the SFP.

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 hop of the SFP at which the metadata is to be introduced. Classifiers and reclassifiers will know the correct SI values to used from information supplied by the control or management plane as is the case for NSH packets with payload data.

5.2. Per-Flow Metadata

Per-flow metadata is metadata that applies to a subset of the packets on an SFP, such as packets matching a particular 5 tuple of source address, destination address, source port, destination port, and payload protocol. This metadata also does not need to be transmitted with every packet, but can be installed at the SFIs on the SFP and applied to the packets that match the flow description.

If there is just one flow on an SFP then there is no difference between per-flow metadata and per-SFP metadata as described in .

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. However, when "Next Protocol" indicates "None" this cannot be done. In this case the identity of the flow is carried in the metadata.

5.3. Coordination Between SFC-Aware SFIs

A pair of SFC-aware SFIs (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 SFIs 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 unnecessary 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 SFI so an implementation of an SF may also be able to generate NSH packets as described in this document.

Note also that a system with SFIs 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. Thus, such an SFI does not need to insert NSH packets onto SFPs used to carry payload data, but can use (and know the SPI of) this special, dedicated SFP.

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 predictable intervals.

- Measuring path qualities distinct from behavior of SFIs.
- 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 specifically 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). Further details of this process are out of scope of this document, but it should be understood that, just as for OAM, an essential feature of using a control channel is that the various speakers are assigned identifiers (i.e., addresses). In this case, those identifiers could be SPI/SI pairs, or could be IP addresses as used in the normal control and management plane of the SFC network.

5.6. Non-Applicable Use Cases

Per-packt metadata is metadata that applies specifically to a payload packet. It informs an SFI how to handle the payload packet, and does not apply to any other packets.

The mechanisms described in this document are not applicable to per-packet metadata because, by definition, if the "Next Protocol" indicates "None" then there is no packet following the NSH for the metadata to be associated with.

6. Security Considerations

Metadata-only packets as enabled by this document provide a covert channel. However, this is only different from the metadata 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 SFIs 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 and forms part of the NSH definition found in [I-D.ietf-sfc-nsh].

The mechanism described 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. The amount of packets with "Next Protocol" set to "None" on an SFP MAY be rate limited at any point on the SFP to provide additional security.

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

Lucy Yong
Retired

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.

Thanks to Eric Rosen and Med Boucadair for valuable review comments.

10. References

10.1. 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>.

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.

[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

Adrian Farrel
Juniper Networks

Email: afarrel@juniper.net

John Drake
Juniper Networks

Email: jdrake@juniper.net