

Service Function Chaining
Internet-Draft
Intended status: Informational
Expires: June 20, 2015

J. Guichard
M. Smith
S. Kumar
Cisco Systems, Inc.
S. Majee
F5 Networks
P. Agarwal
Broadcom
K. Glavin
Riverbed
Y. Laribi
Citrix
December 17, 2014

Network Service Header (NSH) Context Header Allocation (Data Center)
draft-guichard-sfc-nsh-dc-allocation-01

Abstract

This document provides a recommended default allocation for the fixed context headers within a Network Service Header (NSH). NSH is defined in [[I-D.quinn-sfc-nsh](#)]. The allocation scheme is relevant when NSH is used for Service Function Chaining within a data center and may be used to support use cases such as those described in [[I-D.ietf-sfc-dc-use-cases](#)].

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 June 20, 2015.

Copyright Notice

Copyright (c) 2014 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	2
1.1.	Requirements Language	3
2.	Definition Of Terms	3
3.	Network Service Header (NSH) Context Headers	3
4.	Recommended Data Center Mandatory Context Allocation	4
4.1.	Data Center Allocation Specifics	4
5.	Context Allocation and Control Plane Considerations	5
6.	Security Considerations	6
7.	IANA Considerations	6
8.	References	6
8.1.	Normative References	6
8.2.	Informative References	7
	Authors' Addresses	7

[1.](#) Introduction

Network Service Headers (NSH) provide a mechanism to carry shared metadata between network devices and service functions, and between service functions. Such metadata is carried within fixed 32-bit context headers that are part of the NSH structure as defined in [[I-D.quinn-sfc-nsh](#)].

Although NSH also provides the capability to carry variable TLV information following the fixed context headers, the suggested allocation in this draft utilizes the 4 fixed length contexts in order to ensure the broadest possible applicability and support. NSH is carried with packets / frames and is used to create a service plane. The packets / frames are then encapsulated in an outer header for transport.

This document provides a recommended default allocation of these context headers for Service Function Chaining within a data center. The goal of this document is to provide a reference allocation that may be used with or without a control plane. It also serves as a guide to implementers and network operators.

1.1. 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](#)].

2. Definition Of Terms

This document uses the terms as defined in [[I-D.ietf-sfc-problem-statement](#)], [[I-D.quinn-sfc-arch](#)], and [[I-D.quinn-sfc-nsh](#)].

3. Network Service Header (NSH) Context Headers

A Network Service Header in the context of Service Function Chaining is comprised of four parts as described in [[I-D.quinn-sfc-nsh](#)]; a 4-byte base header, a 4-byte service path header, mandatory 4-byte context headers, and optional variable length context headers.

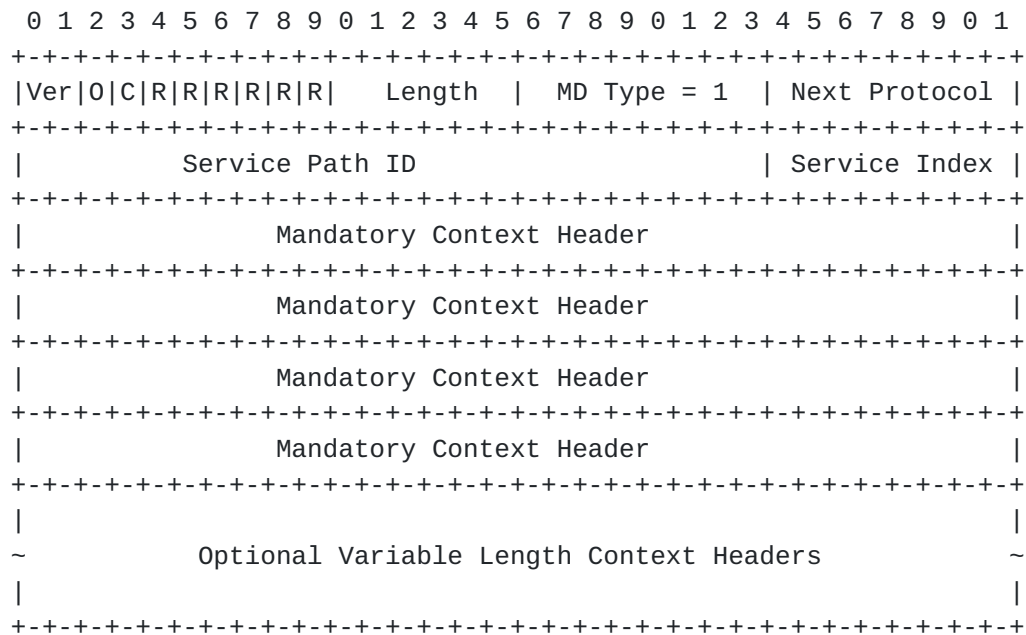


Figure 1: Network Service Header - MD Type 1

4. Recommended Data Center Mandatory Context Allocation

The following context header allocation provides information used to support SFC operation within a generic data center environment.

[[I-D.ietf-sfc-dc-use-cases](#)] provides an overview of data center use cases and requirements to support the allocation.

The 16 bytes of context headers is delivered to service functions that may then use the metadata contained within the headers for local policy enforcement and other functionality.

```

 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|D|F|Res|      Source Node ID      |      Source Interface ID      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Reserved      |      Tenant ID      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Destination Class / Reserved |      Source Class      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|A|      ServiceTag / Opaque Service Class      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 2: NSH DC Context Allocation

4.1. Data Center Allocation Specifics

The specific 16 byte allocation of the mandatory context headers is as follows:

Flag bits: Bits 0-3 are flag bits. Only bits 0 and 1 are defined in this document and the remaining bits are reserved.

D bit: The D-bit is used to indicate whether the Destination Class field in the 3rd word is used. If D-bit is not set then the 3rd word is reserved.

F bit: The F-bit indicates that the 4th word contains a ServiceTag. If F-bit is not set then the 4th word is an opaque field that can be used by service functions.

Source Node ID: An identifier indicating the source device where the original traffic initially entered the Service Function Chain. This identifier is unique within an SFC-enabled domain.

Source Interface ID: An identifier indicating the source interface where the original traffic initially entered the Service Function Chain. This identifier is scoped within the context of the Source Node ID.

Tenant ID: The tenant identifier is used to represent the tenant that the Service Function Chain is being applied to. The Tenant ID is a unique value assigned by a control plane. The distribution of Tenant ID's is outside the scope of this document. As an example application of this field, hardware may insert a VRF ID, VLAN number or VXLAN VNI.

Destination Class: The destination class represents the logical classification of the destination of the traffic. This field is optional and/or the Destination Class may not be known. The D-bit is used to indicate that this field contains a valid Destination Class.

Source Class: represents the logical classification of the source of the traffic. For example, this might represent a source application, a group of like endpoints, or a set of users originating the traffic. This grouping is done for the purposes of applying policy. Policy is applied to groups rather than individual endpoints.

ServiceTag: When the F-bit is set, a ServiceTag is used to identify a particular flow, transaction or an application message unit. The ServiceTag may be used to augment the source and/or destination class. A ServiceTag is a unique identifier that can be used to enable functionality such as classification bypass, slow path skipping and flow programming. As part of the ServiceTag word, bit 0 is the A bit and is used, when needed, to indicate acknowledgment of a ServiceTag by a service function.

5. Context Allocation and Control Plane Considerations

This document describes an allocation scheme for the NSH mandatory context headers defined in [[I-D.quinn-sfc-nsh](#)].

The context header allocations specified in this document are one of many possible allocation schemes and should be used as a guideline only; that is to say these allocations may vary based upon deployment specifics and use cases. The suggested allocation is valid with or without a control plane but the semantics of context values **MUST** be shared amongst participating nodes via some mechanism. The actual method of defining and distributing the allocation scheme is outside of the scope of this document.

6. Security Considerations

This document describes an allocation scheme for the metadata carried within the NSH mandatory context headers. This allocation includes a number of identifiers that must be distributed to participating network elements. While the control plane protocols for distributing these identifiers is outside the scope of this document, any control plane protocol should ensure that these identifiers are securely distributed to the network elements participating in the SFC.

Additionally, many of the fields such as Source and Destination Class described in the metadata directly impact the network policy applied to the traffic flowing through the SFC. There is a risk that these identifiers may be spoofed and proper precautions should be put in place to ensure that these fields can only be updated by trusted entities. Due to the importance of these fields, confidentiality may also be required to ensure that traffic cannot be targeted for attack based on the policy identifiers. This document does not directly address these threats but provides input to the NSH specification as requirements to be considered in securing the contents of the metadata.

7. IANA Considerations

This document includes no request to IANA.

8. References

8.1. Normative References

[I-D.ietf-sfc-dc-use-cases]

Surendra, S., Tufail, M., Majee, S., Captari, C., and S. Homma, "Service Function Chaining Use Cases In Data Centers", [draft-ietf-sfc-dc-use-cases-01](#) (work in progress), July 2014.

[I-D.quinn-sfc-arch]

Quinn, P. and J. Halpern, "Service Function Chaining (SFC) Architecture", [draft-quinn-sfc-arch-05](#) (work in progress), May 2014.

[I-D.quinn-sfc-nsh]

Quinn, P., Guichard, J., Surendra, S., Smith, M., Henderickx, W., Nadeau, T., Agarwal, P., Manur, R., Chauhan, A., Majee, S., Elzur, U., Melman, D., Garg, P., McConnell, B., Wright, C., and K. Kevin, "Network Service Header", [draft-quinn-sfc-nsh-04](#) (work in progress), December 2014.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

8.2. Informative References

[I-D.ietf-sfc-problem-statement]
Quinn, P. and T. Nadeau, "Service Function Chaining Problem Statement", [draft-ietf-sfc-problem-statement-10](#) (work in progress), August 2014.

Authors' Addresses

Jim Guichard
Cisco Systems, Inc.

Email: jguichar@cisco.com

Michael Smith
Cisco Systems, Inc.

Email: michsmit@cisco.com

Surendra Kumar
Cisco Systems, Inc.

Email: smkumar@cisco.com

Sumandra Majee
F5 Networks
90 Rio Robles
San Jose, CA 95134

Email: S.Majee@f5.com

Puneet Agarwal
Broadcom

Email: pagarwal@broadcom.com

Kevin Glavin
Riverbed

Email: Kevin.Glavin@riverbed.com

Youcef Laribi
Citrix

Email: Youcef.Laribi@citrix.com