INTERNET-DRAFT
Intended status: Proposed Standard

Expires: September 4, 2018

Donald Eastlake
Huawei
Bob Briscoe
Independent

March 5, 2018

Network Service Header (NSH) Explicit Congestion Notification (ECN) Support <draft-eastlake-sfc-nsh-ecn-support-00.txt>

Abstract

Explicit congestion notification (ECN) allows a forwarding element to notify downstream devices of the onset of congestion without having to drop packets. This can improve network efficiency through better congestion control without packet drops. This document specifies ECN support within Service Function Chaining (SFC) domains through use of the Network Service Header (NSH).

Status of This Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Distribution of this document is unlimited. Comments should be sent to the SFC Working Group mailing list <sfc@ietf.org>.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/1id-abstracts.html. The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

Table of Contents

1. Introduction

The Network Service Header (NSH [RFC8300]) is used to control the propagation of packets through a Service Function Chaining (SFC [RFC7665]) domain as discussed below. The SFC architecture calls for the encapsulation of traffic inside a service function chaining domain with an NSH being added by the "Classifier" on entry to the domain and the NSH being removed on exit from the domain. Thus the NSH is a natural place to note congestion within the SFC domain, avoiding possible confusion due, for example, to changes in the outer transport header in different parts of the SFC domain.

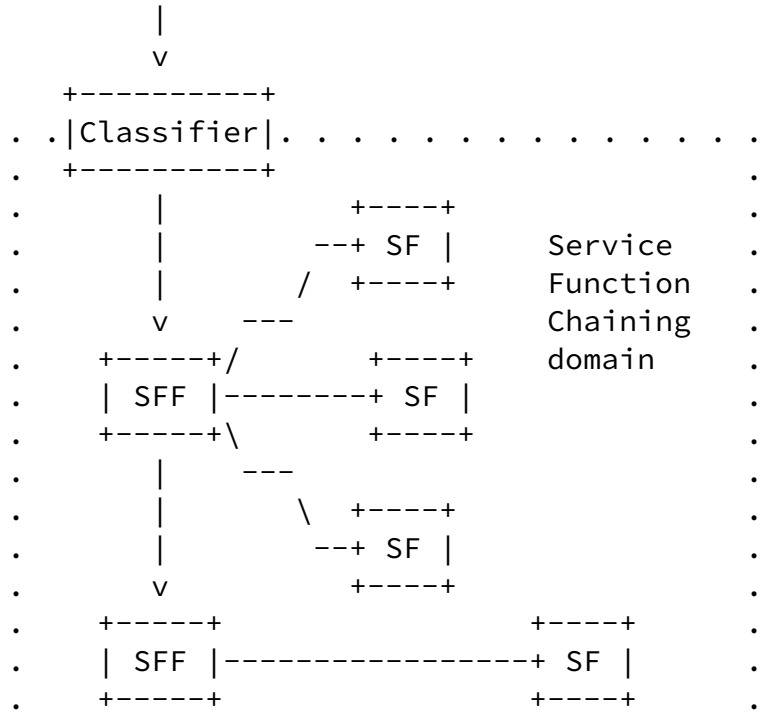

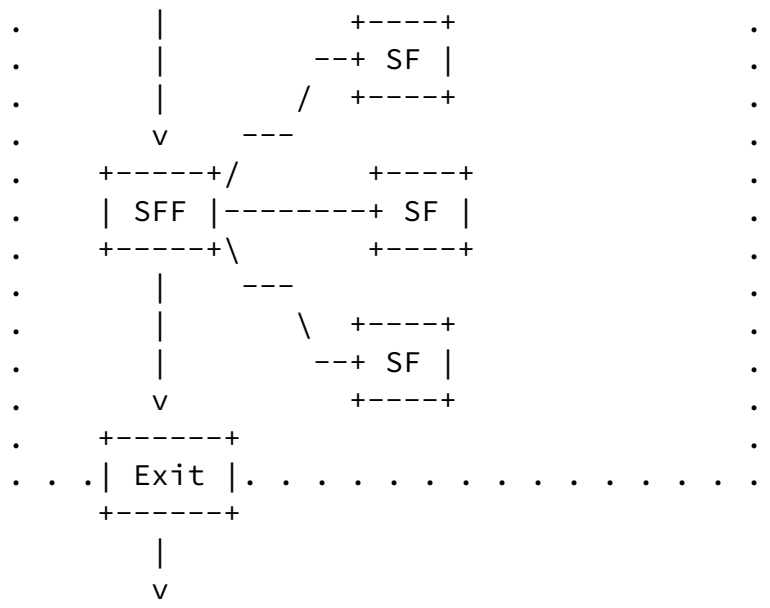

Figure 1. Example Path Forwarding Nodes

Figure 1 shows an SFC domain. Traffic passes through a sequence of Service Function Forwarders (SFFs) each of which sends the traffic to one or more Service Functions (SFs). Each SF performs some operation on the traffic, for example firewall or Network Address Translation (NAT), and returns it to the SFF from which it was received.

Explicit congestion notification (ECN [RFC3168]) allows a forwarding element (such as a router or an SFC Service Function Forwarder (SFF) or Service Function (SF)) to notify downstream devices of the onset of congestion without having to drop packets. This can be used as an element in active queue management or the like [RFC7567] to improve network efficiency through better flow control without packet drops. The forwarding element can explicitly mark a proportion of packets in an ECN field instead of dropping the packet. For example, a two-bit field is available for ECN marking in IP headers [RFC3168].

The availability of congestion information is a building block for various congestion mitigation methods.

1.1 Conventions used in this document

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] [RFC8174] when, and only when, they appear in all capitals, as shown here.

Acronyms:

CE - Congestion Experienced

ECN - Explicit Congestion Notification

ECT - ECN Capable Transport

Not-ECT - Not ECN-Capable Transport

SFC - Service Function Chaining [RFC7665]

2. The NSH ECN Field

Traffic within an SFC domain is encapsulated within an NSH header (see Section 2 of [RFC8300]) as shown in Figure 1.

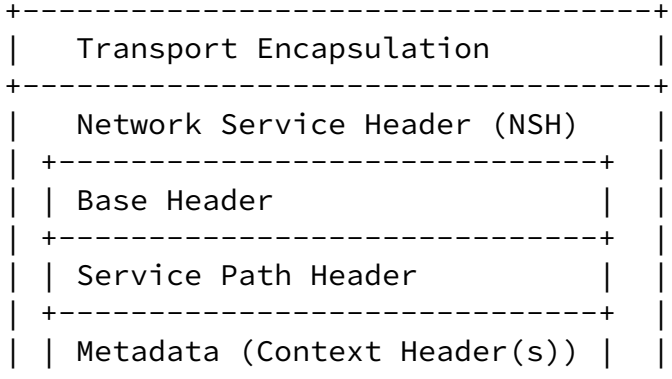

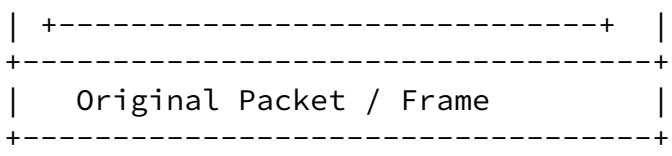

Figure 1. Data Encapsulation in an SFC Domain

Two currently unused bits (indicated by "U") in the NSH Base Header (Section 2.2 of [RFC8300]) are allocated for ECN within the SFC domain as shown in Figure 2.

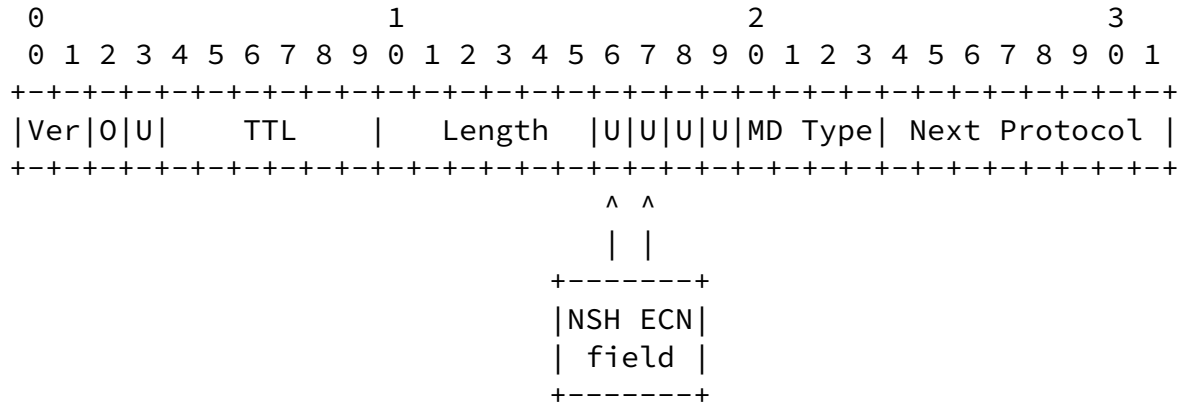

Figure 2: NSH Base Header

Note to RFC Editor: The above figure should be adjusted based on the bits assigned in Section 4 and this note deleted.

Table 1 shows the meaning of the codepoints in the SFC ECN field. These have the same meaning as the ECN field codepoints in the IPv4 or IPv6 header as defined in [RFC3168].

Binary	Name	Meaning
00	Not-ECT	Not ECN-Capable Transport
01	ECT(1)	ECN-Capable Transport
10	ECT(0)	ECN-Capable Transport
11	CE	Congestion Experienced

Table 1. ECN Field Codepoints

3. ECN Support

This section describes the required behavior to support ECN covering the SFC domain.

3.1 At Classifier

When the NSH is added to incoming traffic, the NSH ECN field MUST be set to the ECN-Capable Transport field.

3.2 At SFFs and SFs

Any SFFs and SFs that provides NSH ECN support, if it detects congestion and the NSH ECN field indicates that ECN is supported, MUST set the NSH EC field to the Congestion Experienced value.

3.3 At Exit

In addition to whatever other actions are taken based on Congestion Experienced, if the original packet being carried inside the NSH is IP, the NSH ECN field MUST be combined with IP ECN field as specified in Table 2 that was extracted from [RFC6040]..

Arriving Inner Header	Arriving Outer Header			
	Not-ECT	ECT(0)	ECT(1)	CE
Not-ECT	Not-ECT	Not-ECT	Not-ECT	<drop>
ECT(0)	ECT(0)	ECT(0)	ECT(0)	CE
ECT(1)	ECT(1)	ECT(1)	ECT(1)	CE
CE	CE	CE	CE	CE

Table 2. Exit ECN Fields Merger

4. IANA Considerations

IANA is requested to assign two contiguous bits in the NSH Base Header Bits registry for ECN (bits 16 and 17 suggested) and note this assignment as follows:

Bit	Description	Reference
tbd(16-17)	NSH ECN	[this document]

5. Security Considerations

TBD

For general SFC Security Considerations, see [RFC7665].

6. Acknowledgements

TBD

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>.

[RFC3168] - Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <http://www.rfc-editor.org/info/rfc3168>.

[RFC6040] - Briscoe, B., "Tunnelling of Explicit Congestion Notification", RFC 6040, DOI 10.17487/RFC6040, November 2010, <http://www.rfc-editor.org/info/rfc6040>.

[RFC7567] - Baker, F., Ed., and G. Fairhurst, Ed., "IETF Recommendations Regarding Active Queue Management", BCP 197, RFC 7567, DOI 10.17487/RFC7567, July 2015, <http://www.rfc-editor.org/info/rfc7567>.

[RFC8174] - Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <http://www.rfc-editor.org/info/rfc8174>

[RFC8300] - Quinn, P., Ed., Elzur, U., Ed., and C. Pignataro, Ed., "Network Service Header (NSH)", RFC 8300, DOI 10.17487/RFC8300,

January 2018, <https://www.rfc-editor.org/info/rfc8300>.

Informative References

[RFC7665] - Halpern, J., Ed., and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <https://www.rfc-editor.org/info/rfc7665>.

Authors' Addresses

Donald E. Eastlake, 3rd
Huawei Technologies
155 Beaver Street
Milford, MA 01757 USA

Tel: +1-508-333-2270
Email: d3e3e3@gmail.com

Bob Briscoe
Independent
UK

Email: ietf@bobbriscoe.net
URI: http://bobbriscoe.net/

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. The definitive version of an IETF Document is that published by, or under the auspices of, the IETF. Versions of IETF Documents that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of IETF Documents. The definitive version of these Legal Provisions is that published by, or under the auspices of, the IETF. Versions of these Legal Provisions that are published by third parties, including those that are translated into other languages, should not be considered to be definitive versions of these Legal Provisions. For the avoidance of doubt, each Contributor to the IETF Standards Process licenses each Contribution that he or she makes as part of the IETF Standards Process to the IETF Trust pursuant to the provisions of RFC 5378. No language to the contrary, or terms, conditions or rights that differ from or are inconsistent with the rights and licenses granted under RFC 5378, shall have any effect and shall be null and void, whether published or posted by such Contributor, or included with or in such Contribution.