Explicit Congestion Notification (ECN) and Congestion Feedback
Using the Network Service Header (NSH)
<**draft-eastlake-sfc-nsh-ecn-support-03.txt**>

Abstract

   Explicit congestion notification (ECN) allows a forwarding element to
   notify downstream devices of the onset of congestion without having
   to drop packets. Coupled with a means to feed back information about
   congestion to upstream nodes, this can improve network efficiency
   through better congestion control, frequently without packet drops.
   This document specifies ECN and congestion feedback support within a
   Service Function Chaining (SFC) domain through use of the Network
   Service Header (NSH, RFC 8300) and IP Flow Information Export (IPFIX,
   draft-ietf-tsvwg-tunnel-congestion-feedback).

Status of This Memo

   This Internet-Draft is submitted 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> or to the
   authors.

   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

Explicit congestion notification (ECN [[RFC3168](#)]) allows a forwarding
element to notify downstream devices of the onset of congestion
without having to drop packets. Coupled with a means to feed back
information about congestion to upstream nodes, this can improve
network efficiency through better congestion control, frequently
without packet drops. This document specifies ECN and congestion
feedback support within a Service Function Chaining (SFC [[RFC7665](#)])
domain through use of the Network Service Header (NSH [[RFC8300](#)]) and
IP Flow Information Export (IPFIX [[TunnelCongFeedback](#)]).

It requires that all ingress and egress nodes of the SFC domain
implement ECN. While congestion management will be the most effective
if all interior nodes of the SFC domain implement ECN, some benefit
is obtained even if some interior nodes do not implement ECN. In
particular, congestion at any bottleneck where ECN marking is not
implemented will be unmanaged.

The subsections below in this is section provide background
information on NSH, ECN, congestion feedback, and terminology used in
this document.

[1.1](#) NSH Background

The Service Function Chaining (SFC [[RFC7665](#)]) architecture calls for
the encapsulation of traffic within a service function chaining
domain with a Network Service Header (NSH [[RFC8300](#)]) added by the
"Classifier" (ingress node) on entry to the domain and the NSH being
removed on exit from the domain at the egress node. The NSH is used
to control the path of a packet in an SFC domain. The NSH is a
natural place, in a domain where traffic is NSH encapsulated, to note
congestion, avoiding possible confusion due, for example, to changes
in the outer transport header in different parts of the domain.

```
                 |
                 v
            +----------+
        . .|Classifier|. . . . . . . . . . . .
        . +----------+                          .
        .      |            +----+              .
        .      |          --+ SF |    Service   .
        .      |         /  +----+    Function  .
        .      v      ---           Chaining    .
        .   +-----+/        +----+   domain     .
        .   | SFF |--------+ SF |                .
        .   +-----+\        +----+               .
        .      |    ---                          .
        .      |       \  +----+                 .
        .      |        --+ SF |                 .
        .      v          +----+                 .
        .   +-----+              +----+          .
        .   | SFF |----------------+ SF |        .
        .   +-----+              +----+          .
        .      |           +----+                .
        .      |         --+ SF |                .
        .      |        /  +----+                .
        .      v      ---                        .
        .   +-----+/        +----+               .
        .   | SFF |--------+ SF |                .
        .   +-----+\        +----+               .
        .      |    ---                          .
        .      |       \  +----+                 .
        .      |        --+ SF |                 .
        .      v          +----+                 .
        .   +------+                             .
        . . .| Exit |. . . . . . . . . . . . . .
            +------+
                 |
                 v
```
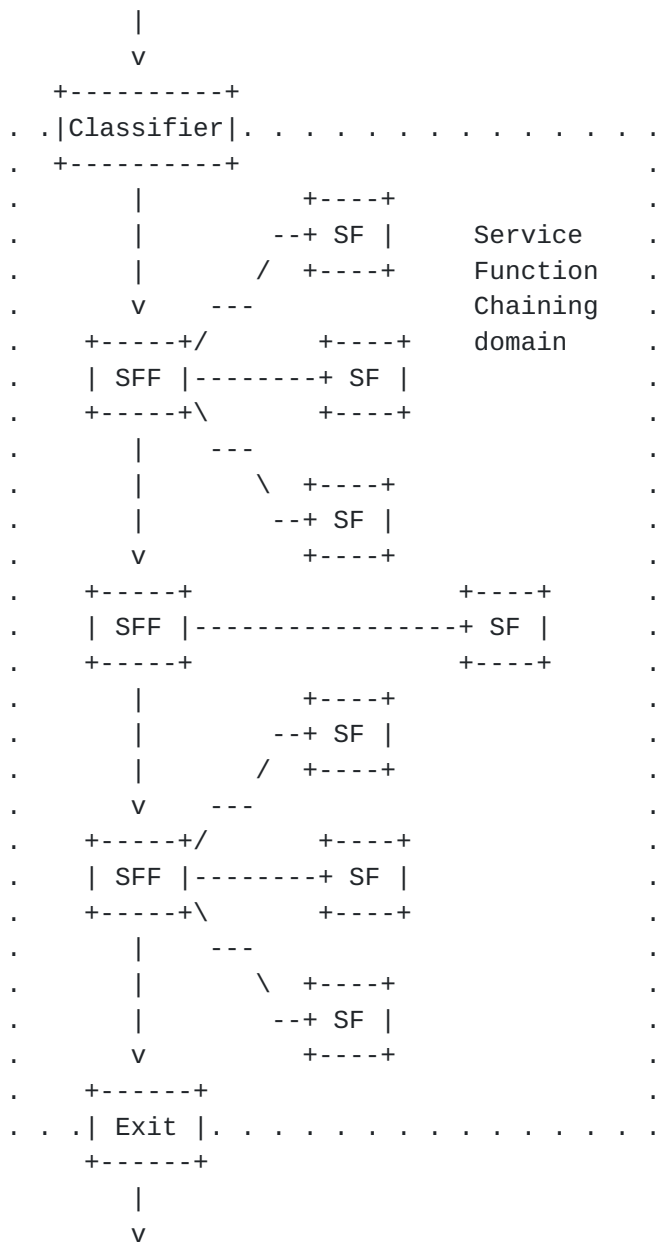
               Figure 1. Example SFC Path Forwarding Nodes

   Figure 1 shows an SFC domain for the purpose of illustrating the use
   of NSH. 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 then returns it to the SFF from which it was received.

   Logically, during the transit of each SFF, the outer transport header
   that got the packet to the SFF is stripped, the SFF decides on the

next forwarding step, either adding a transport header or, if the SFF
is the exit/egress, removing the NSH header. The transport headers

added may be different in different regions of the SFC domain. For example, IP could be used for some SFF-to-SFF communication and MPLS used for other such communication.

1.2 ECN Background

Explicit congestion notification (ECN [RFC3168]) allows a forwarding element (such as a router or an 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 (AQM) [RFC7567] to improve network efficiency through better traffic control without packet drops. The forwarding element can explicitly mark some 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].

1.3 Tunnel Congestion Feedback Background

Tunnel Congestion Feedback [TunnelCongFeedback] is a building block for various congestion mitigation methods. It supports feedback of congestion information from an egress node to an ingress node. Examples of actions that can be taken by an ingress node when it has knowledge of downstream congestion include those listed below. Details of implementing these traffic control methods, beyond those given here, are outside the scope of this document.

Any action by the ingress to reduce congestion needs to allow sufficient time for the end-to-end congestion control loop to respond first, for instance by the ingress taking a smoothed average of the level of congestion signalled by feedback from the tunnel egress.

(1) Traffic throttling (policing), where the downstream traffic flowing out of the ingress node is limited to reduce or eliminate congestion.

(2) Upstream congestion feedback, where the ingress node sends messages upstream to or towards the ultimate traffic source, a function that can throttle traffic generation/transmission.

(3) Traffic re-direction, where the ingress node configures the NSH of some future traffic so that it avoids congested paths. Great care must be taken to avoid (a) significant re-ordering of traffic in flows that it is desirable to keep in order and (b) oscillation/instability in traffic paths due to alternate congestion of previously idle paths and the idling of previously

congested paths. For example, it is preferable to classify

traffic into flows of a sufficiently coarse granularity that the
flows are long lived and use a stable path per flow sending only
newly appearing flows on apparently uncongested paths.

Figure 2 shows an example path from an origin sender to a final
receiver passing through an example chain of service functions
between the ingress and egress of an SFC domain. The path is also
likely to pass through other network nodes outside the SFC domain
(not shown). The figure shows typical congestion feedback that would
be expected from the final receiver to the origin sender, which
controls the load the origin sender applies to all elements on the
path. The figure also shows the congestion feedback from the egress
to the ingress of the SFC domain that is described in this document,
to control or balance load within the SFC domain.

```
  .:= = = = = = = = = = = = = = = = = = = = = = = = = = = = = = :.
 _||_                    End-to-End Congestion Feedback           ||
 \   /                                                            ||
  \/                                                              ||
                     Inner Transport Header and Payload           __
 |  | - - - - - - - - - - - - - - - ->- - - - - - -- - - - - - - |  |
 |  |                                                            |  |
 |  |        .:= = = = = = = = = = = = = = = = = = = = = = =:.   |  |
 |  |      _||_          Tunnel Congestion Feedback        ||    |  |
 |  |      \   /                                           ||    |  |
 |  |       \/                                             ||    |  |
 |  |        __                     NSH                    __    |  |
 |  |       |   |------------------------->--------------|  |    |  |
 |  |. . . |  |       ___          ___          ___      |  | . . .| |
 |  |      |  | | OT1 |    | OT4 |    |   . . . |    | OTn |  |    |  |
 |  |      |  | |-->--|SFF|--->---|SFF|         |SFF|-->--|  |    |  |
 |__|      |__|   |___|         |___|         |___|     |__|    |__|
 origin    SFC     | ^           | ^                     SFC    final
 sender    domain OT2| |OT3   OT6| |OT7                  domain  rcvr
           ingress   v |         v |                     egress
                   +---+       +---+
                   |SF |       |SF |
                   +---+       +---+
```
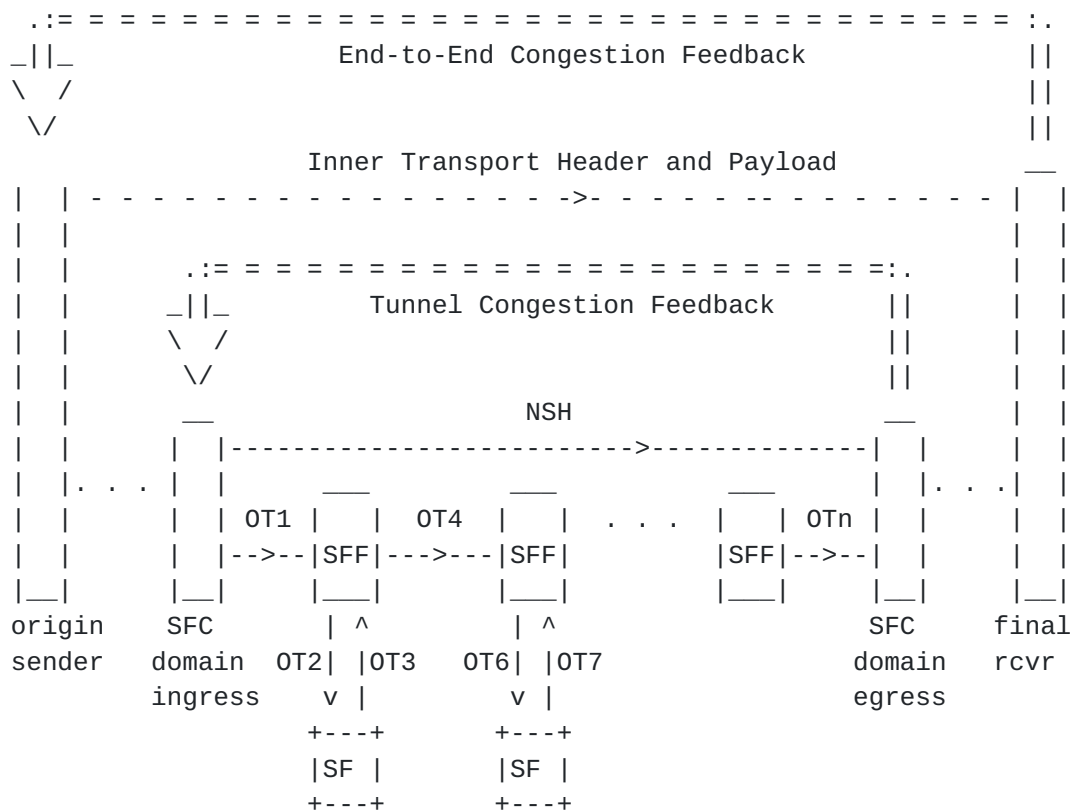
            Figure 2: Congestion Feedback across an SFC Domain

SFC Domain congestion feedback in Figure 2 is shown within the
context of an end-to-end congestion feedback loop. Also shown is the
encapsulated layering of NSH headers within a series of outer
transport headers (OT1, OT2, ... OTn).

1.4 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:

   AQM - Active Queue Management [RFC7567]

   CE - Congestion Experienced [RFC3168]

   downstream - The direction from ingress to egress

   ECN - Explicit Congestion Notification [RFC3168]

   ECT - ECN Capable Transport [RFC3168]

   IPFIX - IP Flow Information Export [RFC7011]

   Not-ECT - Not ECN-Capable Transport [RFC3168]

   NSH - Network Service Header [RFC8300]

   SF - Service Function [RFC7665]

   SFC - Service Function Chaining [RFC7665]

   SFF - Service Function Forwarder [RFC7665] - A type of node that
      forwards based on the NSH.

   TLV - Type Length Value

   upstream - The direction from egress to ingress

[2](). The NSH ECN Field

The NSH header is used to encapsulate and control the subsequent path
of traffic (see [Section 2 of [RFC8300]](#)). The NSH also provides for
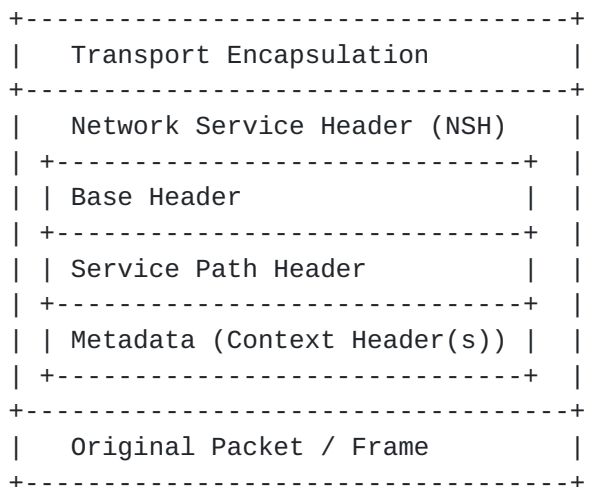metadata inclusion, as shown in Figure 3.

```
               +-----------------------------------+
               |   Transport Encapsulation         |
               +-----------------------------------+
               |    Network Service Header (NSH)    |
               | +-------------------------------+  |
               | | Base Header                   |  |
               | +-------------------------------+  |
               | | Service Path Header           |  |
               | +-------------------------------+  |
               | | Metadata (Context Header(s))  |  |
               | +-------------------------------+  |
               +-----------------------------------+
               |    Original Packet / Frame        |
               +-----------------------------------+
```

Figure 3. Data Encapsulation with the NSH

Two currently unused bits (indicated by "U") in the NSH Base Header
([Section 2.2 of [RFC8300]](#)) are allocated for ECN as shown in Figure
4.

```
 0                   1                   2                   3
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Ver|O|U|     TTL     |   Length  |U|U|U|U|MD Type| Next Protocol |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
                              ^ ^
                              | |
                          +-------+
                          |NSH ECN|
                          | field |
                          +-------+
```
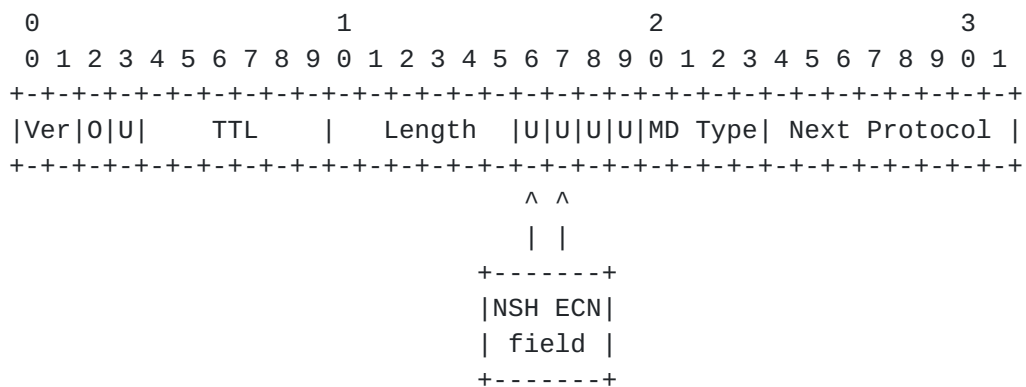
Figure 4: NSH Base Header

Note to RFC Editor: The above figure should be adjusted based on the
bits assigned by IANA (see [Section 5](#)) and this note deleted.

Table 1 shows the meaning of the code points in the NSH ECN field.
These have the same meaning as the ECN field code points 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 Code Points

[3](#). ECN Support in the NSH

   This section describes the required behavior to support ECN using the
   NSH. There are two aspects to ECN support:
   >   1. ECN propagation during encapsulation or decapsulation
   >   2. ECN marking during congestion at bottlenecks.

   While this section covers all combinations of ECN-aware and not ECN-
   aware, it is expected that in most cases the NSH domain will be
   uniform so that, if this document is applicable, all SFFs will
   support ECN; however, some legacy SFs might not support ECN.

   ECN Propagation:

   >   The specification of ECN tunneling [[RFC6040](#)] explains that an
   >   ingress must not propagate ECN support into an encapsulating
   >   header unless the egress supports correct onward propagation of
   >   the ECN field during decapsulation. We define Compliant ECN
   >   Decapsulation here as decapsulation compliant with either
   >   [[RFC6040](#)] or an earlier compatible equivalent ([[RFC4301](#)], or full
   >   functionality mode of [[RFC3168](#)]).

   >   The procedures in [Section 3.2.1](#) ensure that each ingress of the
   >   large number of possible transport links within the SFC domain
   >   does not propagate ECN support into the encapsulating outer
   >   transport header unless the corresponding egress of that link
   >   supports Compliant ECN Decapsulation.

   >   [Section 3.3](#) requires that all the egress nodes of the SFC domain
   >   support Compliant ECN Decapsulation in conjunction with tunnel
   >   congestion feedback, otherwise the scheme in this document will
   >   not work.

   ECN Marking:

   >   At transit nodes the marking behavior specified in 3.2.1 is
   >   recommended and if not implemented at such transit nodes, there
   >   may be unmanaged congestion.

   >   Detection of congestion will be most effective if ECN marking is
   >   supported by all potential bottlenecks inside the domain in which
   >   NSH is being used to route traffic as well as at the ingress and
   >   egress. Nodes that do not support ECN marking, or that support
   >   AQM but not ECN, will naturally use drop to relieve congestion.
   >   The gap in the end-to-end packet sequence will be detected as
   >   congestion by the final receiving endpoint, but not by the NSH
   >   egress (see Figure 2).

3.1 At The Ingress

When the ingress/Classifier encapsulates an incoming IP packet with an NSH, it MUST set the NSH ECN field using the "Normal mode" specified in [RFC6040] (i.e., copied from the incoming IP header).

Then, if the resulting NSH ECN field is Not-ECT, the ingress SHOULD set it to ECT(0). This indicates that, even though the end-to-end transport is not ECN-capable, the egress and ingress of the SFC domain are acting as an ECN-capable transport. This approach will inherently support all known variatns of ECN, including the experimental L4S capability [RFC8311], [ecnL4S].

Packets arriving at the ingress might not use IP. If the protocol of arriving packets supports an ECN field similar to IP, the procedures for IP packets can be used. If arriving packets do not support an ECN field similar to IP, they MUST be treated as if they are Not-ECT IP packets.

Then, as the NSH encapsulated packet is further encapsulated with a transport header, if ECN marking is available for that transport (as it is for IP [RFC3168] and MPLS [RFC5129]), the ECN field of the transport header MUST be set using the "Normal mode" specified in [RFC6040] (i.e., copied from the NSH ECN field).

A summary of these normative steps is given in Table 2.

```
+-----------------+---------------+
| Incoming Header | Departing NSH |
| (also equal to  |   and Outer   |
| departing Inner |    Headers    |
|     Header)     |               |
+-----------------+---------------+
|     Not-ECT     |     ECT(0)    |
|      ECT(0)     |     ECT(0)    |
|      ECT(1)     |     ECT(1)    |
|        CE       |       CE      |
+-----------------+---------------+
```

Table 2. Setting of ECN fields by an ingress/Classifier

The requirements in this section apply to all ingress nodes for the domain in which NSH is being used to route traffic.

3.2 At Transit Nodes

This section described behavior at nodes that forward based on the
NSH such as SFF and other forwarding nodes such as IP routers. Figure
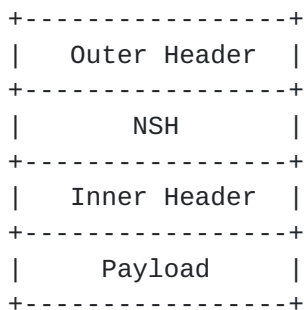5 shows a packet on the wire between forwarding nodes.

```
                  +-----------------+
                  |  Outer Header   |
                  +-----------------+
                  |       NSH       |
                  +-----------------+
                  |  Inner Header   |
                  +-----------------+
                  |     Payload     |
                  +-----------------+
```

Figure 5. Packet in Transit

3.2.1 At NSH Transit Nodes

When a packet is received at an NSH based forwarding node N1, such as
an SFF, the outer transport encapsulation is removed and its ECN
marking SHOULD be combined into the NSH ECN marking as specified in
[RFC6040]. If this is not done, any congestion encountered at non-NSH
transit nodes between N1 and the next upstream NSH based forwarding
node will be lost and not transmitted downstream.

The NSH forwarding node SHOULD use a recognized AQM algorithm
[RFC7567] to detect congestion. If the NSH ECN field indicates ECT,
it will probabilistically set the NSH ECN field to the Congestion
Experienced (CE) value or, in cases of extreme congestion, drop the
packet.

When the NSH encapsulated packet is further encapsulated for
transmission to the next SFF or SF, ECN marking behavior depends on
whether or not the node that will decapsulate the outer header
supports Compliant ECN Decapsulation (see Section 3). If it does,
then the ingress node propagates the NSH ECN field to this outer
encapsulation using the "Normal Mode" of ECN encapsulation [RFC6040]
(it copies the ECN field). If it does not, then the ingress MUST
clear ECN in the outer encapsulation to non-ECT (the "Compatibility
Mode" of [RFC6040]).

3.2.2 At an SF/Proxy

If the SF is NSH and ECN-aware, the processing is essentially the same at the SF as at an SFF as discussed in Section 3.2.1.

If the SF is NSH-aware but not ECN-aware, then the SFF transmitting the packet to the SF will use Compatibility Mode. Congestion encountered in the SFF to SF and SF to SFF paths will be unmanaged.

If the SF is not NSH-aware, then an NSH proxy will be between the SFF and the SF to avoid exposure of the NSH at the SF that does not understand NSHs. This is described in Section 4.6 of [RFC7665]. The SF and proxy together look to the SFF like an NSH-aware SF. The behavior at the proxy and SF in this case is as below:

   If such a proxy is not ECN-aware then congestion in the entire path from SFF to proxy to SF back to proxy to SFF will be unmanaged.

   If the proxy is ECN-aware the proxy uses an AQM to indicate congestion in the proxy itself in the NSH that it returns to the SFF. The outer header used for the proxy to SF path uses Normal Mode. The outer head used for the proxy return to SFF path uses Normal Mode based copying the NSH ECN field to the outer header. Thus congestion in the proxy will be managed. Congestion in the SF will be managed only if the SF is ECN-aware implementing an AQM.

3.2.3 At Other Forwarding Nodes

Other forwarding nodes, that is non-NSH forwarding nodes between NSH forwarding nodes, such as IP routers, might also be potential bottlenecks. If so, they SHOULD implement an AQM algorithm to update the ECN marking in the outer transport header as specified in [RFC3168].

3.3 At Exit/Egress

First, any actions are taken based on Congestion Experienced such as forwarding statistics back to the ingress (see Section 4). If the packet being carried inside the NSH is IP, when the NSH is removed the NSH ECN field MUST be combined with IP ECN field as specified in Table 3 that was extracted from [RFC6040]. This requirement applies to all egress nodes for the domain in which NSH is being used to route traffic.

```
+---------+---------------------------------------------+
|Arriving |           Arriving Outer Header             |
|   Inner +---------+----------+----------+-----------+
|  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 3. Exit ECN Fields Merger

All the egress nodes of the SFC domain MUST support Compliant ECN
Decapsulation as specified in this section. If this is not the case,
the scheme described in this document will not work, and cannot be
used.

3.4 Conservation of Packets

The SFC specification permits an SF to absorb packets and to generate
new packets as well as to process and forward the packets it
receives. Such actions might appear to be packet loss due to
congestion or might mask the loss of packets by generating additional
packets.

The tunnel congestion feedback approach [TunnelCongFeedback] detects
loss by counting payload bytes in at the ingress and counting them
out at the egress. This does not work unless nodes conserve the
amount of payload bytes. Therefore, it will not be possible to detect
loss using this technique if they are not conserved.

Nonetheless, if a bottleneck supports ECN marking, it will be
possible to detect the very high level of CE markings that are
associated with congestion that is so excessive that it leads to
loss. However, it will not be possible for the tunnel congestion
feedback approach to detect any congestion, whether slight or severe,
if it occurs at a bottleneck that does not support ECN marking.

[4](#). **Tunnel Congestion Feedback Support**

   The collection and storage of congestion information may be useful
   for later analysis but, unless it can be fed back to a point which
   can take action to reduce congestion, it will not be useful in real
   time. Such congestion feedback to the ingress enables it to take
   actions such as those listed in [Section 1.3](#).

   IP Flow Information Export (IPFIX [[RFC7011](#)]) provides a standard for
   communicating traffic flow statistics. As extended by
   [[TunnelCongFeedback](#)], IPFIX can be used to determine the extent of
   congestion between an ingress and egress.

   IPFIX recommends use of SCTP [[RFC4960](#)] in partial reliability mode.
   This mode allows loss of some packets, which is tolerable because
   IPFIX communicates cumulative statistics. IPFIX over SCTP SHOULD be
   used directly where there is IP connectivity between the ingress and
   egress; however, there might be different transport protocols or
   address spaces used in different regions of an SFC domain that make
   such direct IP connectivity problematic. The NSH provides the general
   method of routing of traffic within such domain so the IPFIX over
   SCTP over IP traffic should be encapsulated in NSH when necessary.

5. 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]
```

6. **Security Considerations**

   For general NSH security considerations, see [RFC8300].

   For security considerations concerning tampering with ECN signaling,
   see [RFC3168]. For security considerations concerning ECN
   encapsulation, see [RFC6040].

   For general IPFIX security considerations, see [RFC7011]. If deployed
   in an untrusted environment, the signaling traffic between ingress
   and egress can be protected utilizing the security mechanisms
   provided by IPFIX (see section 11 in RFC7011).

   The solution in this document does not introduce any greater
   potential to invade privacy than would have been possible without the
   solution.

7. **Acknowledgements**

   The authors wish to thank the following for their comments and
   suggestion:

      Joel Halpern, Tal Mizrahi, Xinpeng Wei

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>.

   [RFC5129] - Davie, B., Briscoe, B., and J. Tay, "Explicit Congestion
         Marking in MPLS", RFC 5129, DOI 10.17487/RFC5129, January 2008,
         <https://www.rfc-editor.org/info/rfc5129>.

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

   [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, <https://www.rfc-
         editor.org/info/rfc7011>.

   [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>.

   [TunnelCongFeedback] - Wei, X., Zhu, L., and L. Deng, "Tunnel
         Congestion Feedback", draft-ietf-tsvwg-tunnel-congestion-
         feedback, work in progress.

Informative References

   [RFC4301] - Kent, S. and K. Seo, "Security Architecture for the
         Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December
         2005, <https://www.rfc-editor.org/info/rfc4301>.

   [RFC4960] - Stewart, R., Ed., "Stream Control Transmission Protocol",
         RFC 4960, DOI 10.17487/RFC4960, September 2007,
         <https://www.rfc-editor.org/info/rfc4960>.

   [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>.

   [RFC8311] - Black, D., "Relaxing Restrictions on Explicit Congestion
         Notification (ECN) Experimentation", RFC 8311, DOI
         10.17487/RFC8311, January 2018, <https://www.rfc-
         editor.org/info/rfc8311>.

   [ecnL4S] - De Schepper, K., and B. Briscoe, "Identifying Modified
         Explicit Congestion Notification (ECN) Semantics for Ultra-Low
         Queuing Delay (L4S)", draft-ietf-tsvwg-ecn-l4s-id, work in
         progress.

Authors' Addresses

        Donald E. Eastlake, 3rd
        Huawei Technologies
        1424 Pro Shop Court
        Davenport, FL 33896 USA

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

        Bob Briscoe
        Independent
        UK

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

        Andrew G. Malis
        Huawei Technologies

        Email: agmalis@gmail.com