

SFC ECN and Congestion Feedback

draft-ietf-sfc-nsh-ecn-support-00

Donald Eastlake – Huawei d3e3e3@gmail.com

Bob Briscoe – Independent

Andrew Malis – Huawei

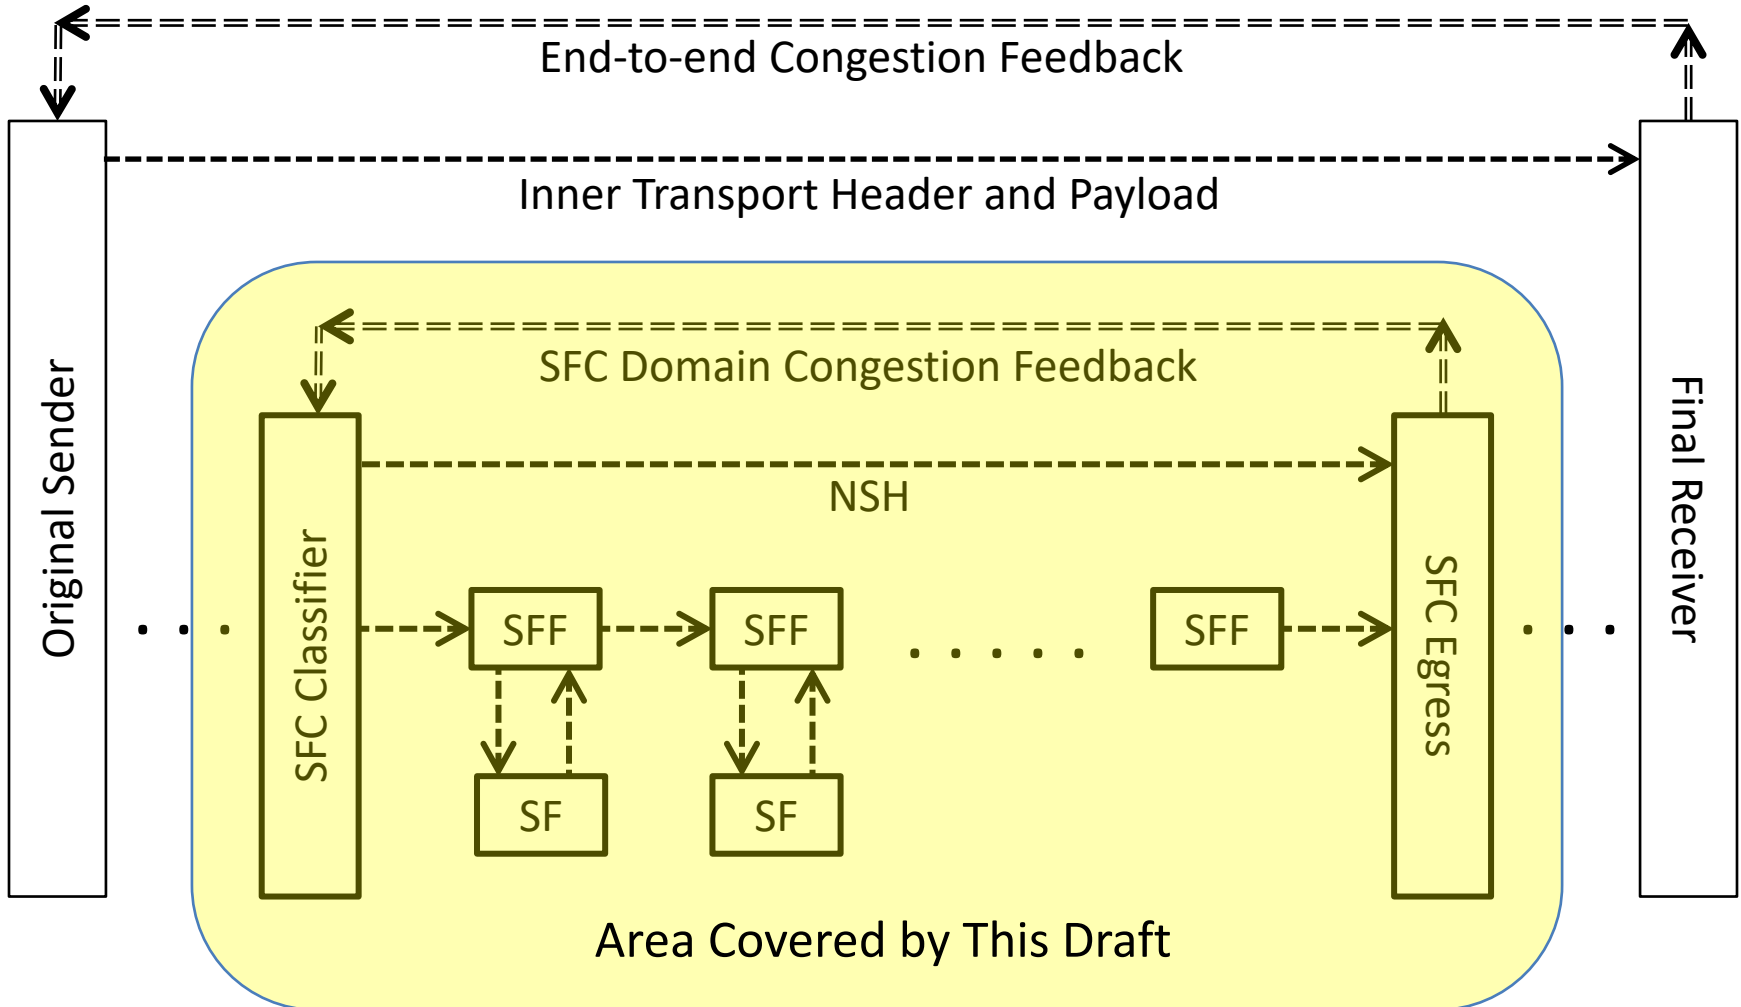
Goals of SFC ECN Support

- Collect congestion information within a Service Function Chaining (SFC) domain with minimal packet drops.
 - Also to contribute to the standardized collection of congestion information from origin before through destination after the SFC domain.
- Communicate SFC domain congestion information to the Classifier(s) so they might take action to reduce congestion.

High Level Overview


- Congestion encountered is communicated downstream towards the SFC domain egress by Explicit Congestion Notification (ECN, RFC 3168) bits in the Network Service Header (NSH, RFC 8300).
- Congestion information is communicated back upstream to the Classifier using IP Flow Information Export (IPFIX, RFC 7011).

High Level Overview



Downstream ECN

- Congestion information is conveyed downstream by two ECN bits in the NSH. Only in severe cases (or where the end-to-end transport does not support ECN) are packets dropped.



Upstream IPFIX

- IPFIX, as extended by draft-ietf-tsvwg-tunnel-congestion-feedback provides mechanisms for communicating, from an egress to a classifier, statistics from which congestions can be determined.
 - Such statistics are cumulative so occasional lost upstream packets are tolerable.
- This is the mandatory mechanism to support of interoperability. Other methods can be used.

Possible Classifier Actions

- Possible classifier actions on learning of congestion include:
 - Traffic throttling.
 - Congestion feedback further upstream.
 - Traffic re-direction.
 - Yes, you have to be very careful to avoid oscillation. For example if you have long lived flows, the Classifier can choose less congested paths but only for newly appearing flows.

Some Details

- This all works better if ECN is implemented throughout the SFC Domain. If an SF does not support NSH or ECN it will have a proxy which should support ECN but even then it is better if the non-NSH supporting SF supports ECN. Basically, any bottleneck where there might be congestion that does not support ECN means that congestion is unmanaged.

Next Steps

- Refine draft, particularly Classified Actions.
- Comments welcome.

- Then request WG Last Call

END

**SFC ECN AND CONGESTION
FEEDBACK**

BACK UP SLIDES

SFC ECN AND CONGESTION FEEDBACK

Some Details

- If the end of an NSH hop (SFF <-> SF, SFF -> SFF) can properly combine Information from the outer transport header into the NSH, then the ingress of such a hop copies the NSH ECN to that transport header. Otherwise, it leaves the outer transport header showing no ECN support and congestion is indicated by packet drop.
 - This requires adding one bit of configuration in each entry at an SFF under a SPI/Index. This bit indicating whether or not the end of the next hop supports ECN when it de-encapsulates.

Some Details

- More graphically

