Network Working Group
Internet-Draft
Intended status: Informational
Expires: November 6, 2015

G. Fairhurst
University of Aberdeen
M. Welzl
University of Oslo

May 05, 2015

The Benefits of using Explicit Congestion Notification (ECN)
draft-ietf-aqm-ecn-benefits-04

Abstract

The goal of this document is to describe the potential benefits when applications use a transport that enables Explicit Congestion Notification (ECN). The document outlines the principal gains in terms of increased throughput, reduced delay and other benefits when ECN is used over network paths that include equipment that supports ECN-marking. It also describes methods that can help successful deployment of ECN. It does not propose new algorithms to use ECN, nor does it describe the details of implementation of ECN in endpoint devices (Internet hosts), routers or other network devices.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Internet Transports (such as TCP and SCTP) are implemented in endpoints (Internet hosts) and have two ways to detect network congestion: the loss of an IP packet or, if Explicit Congestion Notification (ECN) [RFC3168] is enabled, the reception of a packet with a Congestion Experienced (CE)-marking in the IP header. Both of these are treated by transports as indications of congestion. ECN may also be enabled by other transports: UDP applications that provide congestion control may enable ECN when they are able to

correctly process the ECN signals [ID.RFC5405.bis] (e.g., ECN with RTP [RFC6679]).

Active Queue Management (AQM) [ID.RFC2309.bis] is a class of techniques that can be used by network devices (a router, middlebox, or other device that forwards packets through the network) to manage the size of queues in network buffers. A network device that does not support AQM typically uses a drop-tail policy to drop excess IP packets when its queue becomes full. The discard of packets serves as a signal to the end-to-end transport that there may be congestion on the network path being used. This results in a congestion control reaction by the transport to reduce the maximum rate permitted by the sending endpoint.

When an application uses a transport that enables use of ECN [RFC3168], the transport layer sets the ECT(0) or ECT(1) codepoint in the IP header of packets that it sends. This indicates to network devices that they may mark, rather than drop the ECN-capable IP packets. A network device can then signal incipient congestion (network queueing) at a point before a transport experiences congestion loss or additional queuing delay. The marking is generally performed as the result of various AQM algorithms, where the exact combination of AQM/ECN algorithms does not need to be known by the transport endpoints.

Since ECN makes it possible for the network to signal the presence of incipient congestion without incurring packet loss, it lets the network deliver some packets to an application that would otherwise have been dropped if the application or transport did not support ECN. This packet loss reduction is the most obvious benefit of ECN, but it is often relatively modest. However, enabling ECN can also result in a number of beneficial side-effects, some of which may be much more significant than the immediate packet loss reduction from ECN-marking instead of dropping packets. Several benefits reduce latency (e.g., reduced Head-of-Line Blocking). The remainder of this document discusses the potential for ECN to positively benefit an application without making specific assumptions about configuration or implementation.

[RFC3168] describes a method in which a network device sets the CE codepoint of an ECN-Capable packet at the time that the network device would otherwise have dropped the packet. While it has often been assumed that network devices should CE-mark packets at the same level of congestion at which they would otherwise have dropped them, [ID.RFC2309.bis] recommends that network devices allow independent configuration of the settings for AQM dropping and ECN marking. Such separate configuration of the drop and mark policies is supported in some network devices.

The focus of this document is on usage of ECN by transport and application layer flows, not its implementation in endpoint hosts, or in routers and other network devices.

1.1. Terminology

The following terms are used:

Network device: A router, middlebox, or other device that forwards IP packets through the network.

Endpoint: An Internet host that terminates a transport protocol connection across an Intenet path.

non-ECN Capable: An IP packet with a zero value codepoint. A non-ECN capable packet may be forwarded, dropped or queued by a network device.

Incipient Congestion: The detection of congestion when it is starting, perhaps by noting that the arrival rate exceeds the forwarding rate.

CE: Congestion Experienced codepoint, a value marked in the IP packet header.

ECN-Capable: An IP packet with a header marked with a non-zero ECN value (i.e., with a ECT(0), ECT(1), or CE codepoint). An ECN-capable network device may forward, drop or queue a packet and may choose to CE-mark an ECN-capable packet when there is incipient congestion.

2. Benefit of using ECN to avoid Congestion Loss

When a non-ECN capable packet would need to queue or discard as a result of incipient congestion, an ECN-enabled router may be expected to CE-mark, rather than drop an ECN-enabled IP packet [ID.RFC2309.bis]. An application can benefit from this marking in several ways:

2.1. Improved Throughput

ECN can improve the throughput of an application, although this increase in throughput offered by ECN is often not the most significant gain.

When an application uses a light to moderately loaded network path, the number of packets that are dropped due to congestion is small. Using an example from Table 1 of [RFC3649], for a standard TCP sender with a Round Trip Time, RTT, of 0.1 seconds, a packet size of 1500

bytes and an average throughput of 1 Mbps, the average packet drop ratio is 0.02 (i.e., 1 in 50 packets). This translates into an approximate 2% throughput gain if ECN is enabled. In heavy congestion, packet loss may be unavoidable with, or without, ECN.

ECN avoids the inefficiency of dropping data that has already made it across at least part of the network path.

2.2. Reduced Head-of-Line Blocking

Many transports provide in-order delivery of received data segments to the applications they support. When an AQM scheme drops a packet as a signal of incipient congestion, this triggers loss recovery and a congestion control response. For a transport providing in-order delivery, this requires that the transport receiver stalls (or waits) for all data that was sent ahead of a particular segment to be correctly received before it can forward any later data to the application. This is the usual requirement for TCP and SCTP. PR-SCTP [RFC3758], UDP [RFC0768][ID.RFC5405.bis], and DCCP [RFC4340] provide a transport that does not have this requirement. A congestive loss therefore creates a delay of at least one RTT after a loss event before data can be delivered to an application. We call this Head-of-Line (HOL) blocking.

Using ECN, an application continues to receive data when there is incipient congestion. Use of ECN avoids the additional reordering delay in a reliable transport. (When a transport receives a CE-marked packet, it still requests the sender to make an appropriate congestion-response to reduce the maximum transmission rate for future traffic [ID.RFC5405.bis].)

2.3. Reduced Probability of RTO Expiry

A reduction in the possibility of packet loss can be significant for a reliable transport for a class of applications that send a burst of segments and then becomes idle (either because the application has no further data to send or the network prevents sending further data - e.g., flow or congestion control at the transport layer).

Standard transport recovery methods (such as Fast Recovery ([RFC5681]) are often not able to recover from the loss of the last segment (or last few segments) of a traffic burst (also known as a "tail loss"). This is because the receiver is unaware that the lost segments were actually sent, and therefore generates no feedback [Fla13]. Retransmission of these segments therefore relies on expiry of a transport retransmission timer (e.g., expiry of the TCP or SCTP retransmission timeout, RTO [RFC5681]).

A timer expiry results in the consequent loss of state about the network path being used. This typically includes resetting path estimates such as the RTT, re-initialising the congestion window, and possibly updates to other transport state. This can reduce the performance of the transport until it again adapts to the path.

When incipient congestion occurs at the tail of a burst, an ECN-capable network device can CE-mark the packet(s), rather than triggering drop. This allows the transport to avoid the retransmission timeout, which can reduce application level latency and improve the throughput for applications that send intermittent bursts of data and rely upon timer-based recovery of packet loss. The benefit is expected to be especially significant when ECN is used on TCP SYN/ACK packets [RFC5562] where the RTO interval may be large because in this case TCP cannot base the timeout period on prior RTT measurements from the same connection.

2.4. Applications that do not Retransmit Lost Packets

Some latency-critical applications do not retransmit lost packets, yet may be able to adjust the sending rate in the presence of incipient congestion. Examples of such applications include UDP-based services that carry Voice over IP (VoIP), interactive video or real-time data. The performance of many such applications degrades rapidly with increasing packet loss, and many therefore employ mechanisms (e.g., packet forward error correction, data duplication, or media codec error concealment) to mitigate the effect of congestion loss on the application. However, relying on such mechanisms adds complexity and consumes additional network capacity, reducing the available capacity for application data and contributing to the path latency when congestion is experienced.

By decoupling congestion control from loss, ECN can allow the transports supporting these applications to reduce their rate before the application experiences loss from congestion. Because this reduces the negative impact of using loss-hiding mechanisms, ECN can have a direct positive impact on the quality experienced by the users of these applications.

2.5. Making Incipient Congestion Visible

A characteristic of using ECN is that it exposes the presence of congestion on a network path to the transport and network layers. This information can be used for monitoring the level of congestion along the path by a transport/application or a network operator. Metering packet loss is harder. ECN measurements are used by Congestion Exposure (ConEx) [RFC6789].

A network flow that only experiences CE-marking and no loss implies that the sending endpoint is experiencing only congestion and not other sources of packet loss (e.g., link corruption or loss in middleboxes). The converse is not true - a flow may experience a mixture of ECN-marking and loss when there is only congestion, or when there is a combination of packet loss and congestion [ID.RFC2309.bis]. Recording the presence of CE-marked packets can therefore provide information about the current congestion level experienced on a network path. However, it is important to note that any Internet path may also experience congestive loss (e.g., due to queue overflow, AQM methods that protect other flows, middlebox behaviours), so an absence of CE-marks does not indicate a path has not experienced congestion.

2.6. Opportunities for new Transport Mechanisms

Loss is regarded as the standard signal from a network device experiencing congestion. In contrast, CE-marked packets carry an indication that network queues are filling, without incurring loss. This introduces the possibility to provide richer feedback (more frequent and fine-grained indications) to transports. This could utilise new thresholds and algorithms for ECN-marking. ECN therefore provides a mechanism that can benefit evolution of transport protocols.

2.6.1. Benefits from other forms of ECN-Marking/Reactions

ECN requires a definition of both how network devices CE-mark packets and how applications/transports react to reception of these CE-marked packets. ECN-capable receiving endpoints therefore need to provide feedback indicating that CE-marks were received.[RFC3168]provides a method that signals once each round trip time that CE-marked packets have been received. An endpoint may provide more detailed feedback describing the set of received ECN codepoints using Accurate ECN Feedback [ID.Acc.ECN]. This can provide more information to a congestion control mechanism at the sending endpoint.

Loss and CE-marking are both used as an indication for congestion. However, while the amount of feedback that is provided by loss ought naturally to be minimized, this is not the case for ECN. With ECN, a network device could provide richer and more frequent feedback on its congestion state. This could be used by the control mechanisms in the transport to make a more appropriate decision on how to react to congestion, allowing it to react faster to changes in congestion state.

Precise feedback about the number of packet marks encountered is supported by the Real Time Protocol (RTP) when used over UDP [RFC6679] and proposed for SCTP [ST14] and TCP [ID.Acc.ECN].

Benefit has been noted when packets are CE-marked earlier using an instantaneous queue, and if the receiver provides feedback about the number of packet marks encountered, an improved sender behavior has been shown to be possible (e.g, Datacenter TCP (DCTCP) [AL10]). DCTCP is targeted at confined environments such as a datacenter. It is currently unknown whether or how such behaviour could be safely introduced into the Internet.

3. Network Support for ECN

For an application to use ECN requires that the endpoint first enables ECN within the transport.

The ability to use ECN requires network devices along the path to at least forward IP packets with any ECN codepoint (i.e., packets with ECT(0), ECT(1), or with a CE-mark), see also Section 3.3.

For an application to gain benefit from using a transport that enables ECN, network devices need to enable ECN. However, not all network devices along the path need to enable ECN. Any network device that does not CE-mark an ECN-enabled packet can be expected to drop packets under congestion. Applications that experience congestion in these network devices do not see any benefit from using ECN, but would see benefit if the congestion were to occur within a network device that did support ECN.

There is opportunity to design an AQM method for ECN that differs from one designed to drop packets (e.g., Random Early Detection uses a smoothed queue length because it was designed for loss and a congestion control that halves its sending rate on congestion) [ID.RFC2309.bis]. IETF-specified AQM algorithms also need to be designed to work with network paths that may contain multiple bottlenecks. Transports can therefore experience dropped or CE-marked packets from more than one network device related to the same network flow [ID.AQM.eval].

ECN can be deployed both in the general Internet and in controlled environments:

- o ECN can be incrementally deployed in the general Internet. The IETF has provided guidance on configuration and usage in [ID.RFC2309.bis]. A recent survey reported a growing support for network paths to pass ECN codepoints [TR15].

- ECN may also be deployed within a controlled environment, for example within a data centre or within a well-managed private network. In this case, the use of ECN may be tuned to the specific use-case. An example is Datacenter TCP (DCTCP) [AL10] [ID.DCTCP].

Some mechanisms can assist in using ECN across a path that only partially supports ECN. These are noted in Section 4. Applications and transports (such as TCP or SCTP) can be designed to fall-back to not using ECN when they discover they are using a path that does not allow use of ECN (e.g., a firewall or other network device configured to drop the ECN codepoint) Section 4.2.

3.1. Enabling ECN in Network Devices

All network devices need to be configured not to drop packets solely because the ECT(0) or ECT(1) codepoints are used.

The ECN behaviour of a network device should be configurable [ID.RFC2309.bis]. An AQM algorithm that supports ECN needs to define the threshold and algorithm for ECN-marking.

A network device must not set the CE-mark in a packet, except to signal incipient congestion, since this will be interpreted as incipient congestion by the transport endpoints.

3.2. Tunneling ECN and the use of ECN by Lower Layer Networks

Some networks may use ECN internally or tunnel ECN (e.g., for traffic engineering or security). Guidance on the correct use of ECN in this case is provided in [RFC6040].

Further guidance on the encapsulation and use of ECN by non-IP network devices is provided in [ID.ECN-Encap].

3.3. Bleaching and Middlebox Requirements to deploy ECN

Cases have been noted where a sending endpoint marks a packet with a non-zero ECN mark, but the packet is received with a zero ECN codepoint by the remote endpoint [TR15]. This could be a result of a policy that erases or "bleaches" the ECN codepoint values at a network edge (resetting the codepoint to zero).

Bleaching may occur for various reasons (including normalising packets to hide which equipment supports ECN). The current IPv4 and IPv6 specifications assign usage of 2 bits in the IP header to carry the ECN codepoint. This 2-bit field was reserved in [RFC2474] and assigned in [RFC3168]. A previous usage assigned these bits as a

part of the now deprecated Type of Service (ToS) field [RFC1349]. A network device that conforms to this older specification may remark or erase the ECN codepoints, and such equipment needs to be updated to the current specifications to support ECN.

This policy prevents use of ECN by applications. A network device should therefore not remark an ECT(0) or ECT(1) mark to zero. This can result in IP packets that were originally marked as ECN-capable being dropped by ECN-capable network devices further along the path, and eliminates the advantage of using of ECN.

A network device must not change a packet with a CE mark to a zero codepoint (if the CE marking is not propagated, a CE-marked packet must be discarded) [ID.RFC2309.bis]. A CE-marked packet should be expected to have already received ECN treatment in the network, and remarking it would then hide the congestion signal from the receiving endpoint. This eliminates the benefits of ECN. It can also slow down the response to congestion compared to using AQM, because the transport will only react if it later discovers congestion by some other mechanism. (Note that ECN-enabled network devices need to drop excessive traffic [ID.RFC2309.bis], even when this is marked as ECN-capable.)

3.4. Impact on non-ECN flows

A simple AQM scheme could place ECN-capable and non-ECN capable flows withing the same queue. Since an ECN AQM scheme would normally CE-mark packets during incipient congestion, these packets would not be removed from a queue, in contrast to discarding the IP packet in a drop-based AQM scheme. Design of an appropriate queue management system needs to therefore consider when packets are dropped due to incipient congestion, and seek to provide appropriate fairness between ECN and non-ECN traffic, e.g. using more advanced AQM methods or including flow isolation using scheduling [ID.RFC2309.bis].

4. Using ECN across the Internet

This section describes partial deployment of ECN.

Early use of ECN by transports/applications encountered a number of operational difficulties when the network path either failed to transfer ECN-capable packets or the remote endpoint did not fully support use of ECN. The remainder of the section describes transport mechanisms that allow ECN-enabled endpoints to continue to work effectively over a path with misbehaving network devices or to detect and react to non-conformant paths.

4.1. Partial Deployment

ECN has been designed to allow incremental partial deployment [RFC3168]. Any network device may choose to use either ECN or some other loss-based policy to manage its traffic. Similarly, negotiation allows senders and receivers at the transport layer to choose whether ECN is to be used to manage congestion for a particular network flow.

Usability problems were reported in early deployment of ECN and have been observed to diminish with time, but may still be encountered on some Internet paths [TR15].

4.2. Verifying whether a Path Really Supports ECN

ECN transport and applications need to implement mechanisms to verify ECN support on the entire path that they use (including the remote endpoint) and fall back to not using ECN when it would not work. This is expected to be a normal feature of IETF-defined transports supporting ECN.

Before a transport relies on the presence or absence of CE-marked packets, it may need to verify that any ECN marks applied to packets passed by the path are indeed delivered to the remote endpoint. This could be achieved by the sender setting known ECN codepoints into specific packets in a network flow and then verifying that these reach the remote endpoint [ID.Fallback], [TR15].

The design of any transport/application also needs to be robust to path changes. A change in the set of network devices along a path could impact the ability to effectively signal or use ECN across the path, e.g., when a path changes to use a middlebox that bleaches ECN codepoints. As a necessary, but short term fix, transports could implement mechanisms that detect this and fall-back to disabling use of ECN [BA11].

4.3. Detecting ECN Receiver Feedback Cheating

It is important that receiving endpoints accurately report the loss they experience when using a transport that uses loss-based congestion control. So also, when using ECN, a receiver must correctly report the congestion marking that it receives by providing a mechanism to feed this congestion information back to the sending endpoint.

The transport at endpoint receivers must not try to conceal reception of CE-marked packets in the ECN feedback information that they provide to the sending endpoint [ID.RFC2309.bis]. Transport

protocols are actively encouraged to include mechanisms that can detect and appropriately respond to such misbehavior (e.g., disabling use of ECN, and relying on loss-based congestion detection [TR15]).

5. Summary: Enabling ECN in Network Devices and Hosts

This section provides a list of key requirements to achieve ECN deployment. It also summarises the benefits of deploying and using ECN within the Internet.

Network devices should enable ECN and people configuring host stacks should also enable ECN [ID.RFC2309.bis].

Prerequisites for network devices (including IP routers) to enable use of ECN include:

- must not change a packet with a CE mark to a zero codepoint (if the CE marking is not propagated, the packet must be discarded).

- should not reset the ECN codepoint to zero by default (see Section 3.3).

- should correctly update the ECN codepoint in the presence of congestion [ID.RFC2309.bis].

- may support alternate ECN semantics [RFC4774].

Prerequisites for network endpoints to enable use of ECN include:

- should use transports that can set and receive ECN marks.

- when ECN is used, must correctly return feedback indicating congestion to the sending endpoint.

- when ECN is used, must use transports that react appropriately to received ECN feedback (see Section 4.3).

- when ECN is used, should use transports that can detect misuse of ECN and detect paths that bleach ECN, providing fallback to loss-based congestion detection when ECN is not supported (see Section 4.2).

Application developers should where possible use transports that enable the benefits of ECN. Applications that directly use UDP need to provide support to implement the functions required for ECN [ID.RFC5405.bis]. Once enabled, an application that uses a transport that supports ECN will experience the benefits of ECN as network deployment starts to enable ECN. The application does not need to be

rewritten to gain these benefits. Table 1 summarises some of these benefits.

Section	Benefit
2.1	Improved throughput
2.2	Reduced Head-of-Line blocking
2.3	Reduced probability of RTO Expiry
2.4	Applications that do not retransmit lost packets
2.5	Making incipient congestion visible
2.6	Opportunities for new transport mechanisms

Table 1: Summary of Key Benefits

6. Acknowledgements

The authors were part-funded by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700). The views expressed are solely those of the authors.

The authors would like to thank the following people for their comments on prior versions of this document: Bob Briscoe, David Collier-Brown, John Leslie, Colin Perkins, Richard Scheffenegger, Dave Taht, Wes Eddy, Fred Baker, Mikael Abrahamsson, Mirja Kuehlewind, John Leslie, and other members of the TSVWG.

7. IANA Considerations

XX RFC ED - PLEASE REMOVE THIS SECTION XXX

This memo includes no request to IANA.

8. Security Considerations

This document introduces no new security considerations. Each RFC listed in this document discusses the security considerations of the specification it contains.

9. Revision Information

XXX RFC-Ed please remove this section prior to publication.

Revision 00 was the first WG draft.

Revision 01 includes updates to complete all the sections and a rewrite to improve readability. Added section 2. Author list reversed, since Gorry has become the lead author. Corrections following feedback from Wes Eddy upon review of an interim version of this draft.

Note: Wes Eddy raised a question about whether discussion of the ECN Pitfalls could be improved or restructured - this is expected to be addressed in the next revision.

Revision 02 updates the title, and also the description of mechanisms that help with partial ECN support.

We think this draft is ready for wider review. Comments are welcome to the authors or via the IETF AQM or TSVWG mailing lists.

Revision 03 includes updates from the mailing list and WG discussions at the Dallas IETF meeting.

The section "Avoiding Capacity Overshoot" was removed, since this refers primarily to an AQM benefit, and the additional benefits of ECN are already stated. Separated normative and infoirmative references

Revision 4 (WG Review during WGLC)

Updated the abstract.

Added a table of contents.

Addressed various (some conflicting) comments during WGLC with new text.

The section on Network Support for ECN was moved, and some suggestions for rewording sections were implemented.

Decided not to remove section headers for 2.1 and 2.2 - to ensure the document clearly calls-out the benefits.

Updated references. Updated text to improve consistency of terms and added deifinitions for key terms.

Note: The group suggested this document should not define recommendations for end hosts or routers, but simply state the things needs to enable deployment to be sucessful.

10. References

10.1. Normative References

[ID.RFC2309.bis]
Baker, F. and G. Fairhurst, "IETF Recommendations Regarding Active Queue Management", Internet-draft draft-ietf-aqm-recommendation-06, October 2014.

[ID.RFC5405.bis]
Eggert, Lars., Fairhurst, Gorry., and Greg. Shepherd, "Unicast UDP Usage Guidelines", 2015.

[RFC2474] "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers".

[RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, September 2001.

[RFC6040] Briscoe, B., "Tunnelling of Explicit Congestion Notification", RFC 6040, November 2010.

10.2. Informative References

[AL10] Alizadeh, M., Greenberg, A., Maltz, D., Padhye, J., Patel, P., Prabhakar, B., Sengupta, S., and M. Sridharan, "Data Center TCP (DCTCP)", SIGCOMM 2010, August 2010.

[BA11] Bauer, Steven., Beverly, Robert., and Arthur. Berger, "Measuring the State of ECN Readiness in Servers, Clients, and Routers, ACM IMC", 2011.

[Fla13] Flach, Tobias., Dukkipati, Nandita., Terzis, Andreas., Raghavan, Barath., Cardwell, Neal., Cheng, Yuchung., Jain, Ankur., Hao, Shuai., Katz-Bassett, Ethan., and Ramesh. Govindan, "Reducing web latency: the virtue of gentle aggression.", SIGCOMM 2013, October 2013.

[ID.AQM.eval]
Kuhn, Nicolas., Natarajan, Preethi., Ros, David., and Naeem. Khademi, "AQM Characterization Guidelines (Work-in-progress, draft-ietf-aqm-eval-guidelines)", 2015.

[ID.Acc.ECN]
Briscoe, Bob., Scheffeneger, Richard., and Mirja. Kuehlewind, "More Accurate ECN Feedback in TCP, Work-in-Progress".

[ID.DCTCP]
Bensley, S., Eggert, Lars., and D. Thaler, "Microsoft's Datacenter TCP (DCTCP): TCP Congestion Control for Datacenters (Work-in-progress, draft-bensley-tcpm-dctcp)", 2015.

[ID.ECN-Encap]
Briscoe, B., Kaippallimalil, J., and P. Thaler, "Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP", Internet-draft, IETF work-in-progress draft-ietf-tsvwg-ecn-encap-guidelines.

[ID.Fallback]
Kuehlewind, Mirja. and Brian. Trammell, "A Mechanism for ECN Path Probing and Fallback, draft-kuehlewind-tcpm-ecn-fallback, Work-in-Progress".

[RFC0768] Postel, J., "User Datagram Protocol", 1980.

[RFC1349] "Type of Service in the Internet Protocol Suite".

[RFC3649] Floyd, S., "HighSpeed TCP for Large Congestion Windows", RFC 3649, December 2003.

[RFC3758] Stewart, R., Ramalho, M., Xie, Q., Tuexen, M., and P. Conrad, "Stream Control Transmission Protocol (SCTP) Partial Reliability Extension", RFC 3758, May 2004.

[RFC4340] Kohler, E., Handley, M., and S. Floyd, "Datagram Congestion Control Protocol (DCCP)", RFC 4340, March 2006.

[RFC4774] Floyd, S., "Specifying Alternate Semantics for the Explicit Congestion Notification (ECN) Field", BCP 124, RFC 4774, November 2006.

[RFC5562] Kuzmanovic, A., Mondal, A., Floyd, S., and K. Ramakrishnan, "Adding Explicit Congestion Notification (ECN) Capability to TCP's SYN/ACK Packets", RFC 5562, June 2009.

[RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", RFC 5681, September 2009.

[RFC6679] Westerlund, M., Johansson, I., Perkins, C., O'Hanlon, P., and K. Carlberg, "Explicit Congestion Notification (ECN) for RTP over UDP", RFC 6679, August 2012.

[RFC6789] Briscoe, B., Woundy, R., and A. Cooper, "Congestion Exposure (ConEx) Concepts and Use Cases", RFC 6789, December 2012.

[ST14] Stewart, R., Tuexen, M., and X. Dong, "ECN for Stream Control Transmission Protocol (SCTP)", Internet-draft draft-stewart-tsvwg-sctpecn-05.txt, January 2014.

[TR15] Tranmmel, Brian., Kuehlewind, Mirja., Boppart, Damiano, Learmonth, Iain., and Gorry. Fairhurst, "Enabling internet-wide deployment of Explicit Congestion Notification Tramwell, B., Kuehlewind, M., Boppart, D., Learmonth, I., Fairhurst, G. & Scheffnegger, Passive and Active Measurement Conference (PAM)", March 2015.

Authors' Addresses

Godred Fairhurst
University of Aberdeen
School of Engineering, Fraser Noble Building
Aberdeen AB24 3UE
UK

Email: gorry@erg.abdn.ac.uk

Michael Welzl
University of Oslo
PO Box 1080 Blindern
Oslo N-0316
Norway

Phone: +47 22 85 24 20
Email: michawe@ifi.uio.no