

TSVWG  
Internet-Draft  
Intended Status: Informational  
Expires: April 4, 2013

K. Carlberg  
G11  
P. O'Hanlon  
UCL  
Oct 4, 2012

**Reactions to Signaling from ECN Support for RTP/RTCP**  
**<[draft-carlberg-tsvwg-ecn-reactions-03.txt](#)>**

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."

Copyright Notice

Copyright (c) 2012 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.



## Abstract

This document presents various responses to Congestion Experience (CE) notifications by real time applications that have negotiated end-to-end support of Explicit Congestion Notification (ECN). This document is a follow-on effort of [[rfc6679](#)], which specifies the signaling used to provide ECN support for RTP/RTCP flows.

## **1. Introduction**

This document presents various responses to Congestion Experience (CE) notifications by real time applications that have negotiated end-to-end support of Explicit Congestion Notification (ECN). [[rfc6679](#)] defines the signaling for support of ECN by RTP based sessions, and also covers the case where a set of nodes do not respond to CE notifications. A more detailed discussion about how back-off algorithms can be achieved and supported for specific applications is viewed as out of scope of that document and may be addressed by a companion document.

### **1.1 Background**

ECN is a mechanism used to explicitly signal the presence of congestion without relying on packet loss. It was initially designed using a dual layer signaling model; negotiation and feedback at the transport layer, and downstream notification of congestion at the network layer. For IP, a new two bit field was used to both indicate the successful negotiated support for ECN signaling, as well as indicate the presence of congestion via the CE flag. In the case of TCP [[rfc3168](#)], a new TCP header flag was defined that provides upstream end-to-end indication of congestion occurring somewhere along the downstream path.

There should be no difference in congestion response if ECN-CE marks or packet drops are detected. However it is noted that there MAY be other reactions to ECN-CE specified in the future. Such an alternative reaction MUST be specified and considered to be safe for deployment under any restrictions specified. We specify such an alternative in this document.

With respect to ECN for TCP, [[rfc3168](#)] specifies an indication of congestion, but it does so once per Round Trip Time (RTT). [[rfc6679](#)] is an effort that proposes a finer grained notification reflecting a more accurate indication of the number of ECN marked packets received within one RTT.

### **1.2 Terminology and Abbreviations**

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](#) [[RFC2119](#)].

## **2. Issues**

The initial discussions and presentation of [[rfc6679](#)] produced a consensus that the specification of signaling was to be done within the AVTcore working group, and any subsequent discussion on end-to-end reactions to the signaling would be accomplished in the Transport Services (TSV) working group. This draft satisfies the latter effort.

Another issue that needs to be recognized is that the reactions to CE in the context of [[rfc6679](#)] are the responsibility of the application. This is in contrast to ECN support for TCP, where explicit signaled feedback of, and reaction to, CE is kept transparent to the application. The issue of placing the feedback responsibility in the application is that each application needs to add specific support for that reaction. On the other hand, multiple reactions may be considered by the application. For this reason, [[rfc6679](#)] states the need for a default congestion control reaction that MUST be supported. [Section 3](#) through 5 expands on this topic.

## **3. Congestion Control Algorithms**

The transport of any data flow across the Internet produces a need for some form of congestion control to attain a suitable share of the capacity of the path through a network. Most of the existing work on realtime congestion control algorithms has been rooted in TCP-friendly approaches but with smoother adaptation cycles. TCP congestion control is unsuitable for interactive media for a number of reasons including the fact that it is loss-based so it maximises the latency on a path, it changes its transmit rate to quickly for multimedia, and favours reliability over timeliness. In the case of real time media transport, one requires:

Smoother rate variation: (than for bulk data) to accommodate the underlying media flow's characteristics.

Low latency: Maintaining latencies sufficient to be usable, where 150ms one way delay is understood to be a good target [[ITU.G114.2003](#)].

Fairness: The algorithm must be fair to both itself and other flows

### **[3.1](#) TCP Friendly Rate Control (TFRC)**

TFRC has a smoother response to congestion than TCP-like approaches, thus making it more suitable for real-time interactive multimedia



applications. It has been cited in a number of other documents within the IETF for use with UDP and media flows [[rfc3714](#), [bcp145](#)] and is seeing full and partial deployment in related solutions such as Empathy/Farsight, and GoogleTalk [[goog1](#)].

However it should be noted that TFRC is only recommended for real-time media use with ECN response. TFRC is not recommended for non-ECN paths due to its loss based operation which leads to full queues with maximised latencies. It is assumed that ECN markings will usually occur with lower queue occupancy and thus lower latency. However it is understood that ECN marks may not provide for sufficiently low latencies in some situations so other congestion control solutions would be preferable.

[[rfc4342](#)] specifies the profile for TFRC for use in the Datagram Congestion Control Protocol (DCCP) [[rfc4340](#)] for a half connection. A DCCP half connection is defined as application data sent downstream with corresponding acknowledgements sent upstream. These half-connections can be realized in the form of one-way pre-recorded media, one-way live media, or two-way interactive. A perceived drawback in this profile concerns its application to interactive media that use small packets. [[RFC4828](#)] is an experimental protocol defining a variation of TFRC used to address this drawback and achieve the same bandwidth as a TCP flow using packets of size 1500 bytes.

[[rfc6679](#)] is a standard that specifies how RTP flows can be supported using the RTP/AVPF profile and the general RTP header extension mechanism.

## **[3.2](#) Related Work**

### **[3.2.1](#) 3GPP**

Outside of this previous and on-going work with TFRC, it is understood that some parties have issues with the behavior of TFRC under certain conditions. A notable mention of this is made in the 3GPP's document on IP Multimedia Subsystem (IMS) Media handling and interaction [[TR26.114](#)], where it is mentioned:

"Note that for IMS networks, which normally have nonzero packet loss and fairly long round-trip delay, the amount of bitrate reduction specified in [RFC 3448](#) is generally too restrictive for video and may, if used as specified, result in very low video bitrates already at (for IMS) moderate packet loss rates."

Though it is unclear exactly what the 3GPP community consider as too restrictive and whether some alteration of the response may be suitable. It should be noted that the 3GPP document only referred to an older





version of TFRC defined in [[RFC3448](#)]. Given that the current version of TFRC [[RFC5348](#)] has made significant changes to the idle and data-limited responses it is unclear whether their assessment is relevant to current TFRC implementations.

Furthermore the specification [TR26.114] only outlines a rudimentary approach to congestion control, providing an example of a 60% back-off reaction to loss within an RTCP reporting period. The proposed signalling employs Temporary Maximum Media Stream Bit Rate Request (TMMBR) [[RFC5104](#)] and Codec Mode Request (CMR) [[RFC4867](#)] for video and audio respectively, which would only provide for very basic rate control if used as specified. We note that [TR26.114] specifies terminal behavior, while [TS36.300] specifies base station behaviour, though neither specify any standardised congestion control approach.

It is understood that there are a number of proprietary and patented approaches that provide more sophisticated response in the case of 3G/LTE, but since these are neither endorsed nor standardized this document advocates a standardized approach such as TFRC.

We also acknowledge that there are many congestion control algorithms available for implementers to choose from, with a subset that are specifically suited to real time media transmission. However, given a variety of real time applications and their various characteristics (sender-only broadcast, interactive unicast, etc), we need to expand the notion of how back-off can be achieved. Hence, the focus needs to be on an output that would resemble the characteristics of TFRC.

Within the RTCweb Working Group the need for a more media friendly congestion control mechanism has been made apparent. Currently, TFRC is perceived as having deficiencies (e.g. its loss-based design, lack of cross-stream congestion control functionality etc) that make it an incomplete or insufficient solution for the envisioned RTCWEB media flows. The RTP Media Congestion Avoidance Techniques (rmcat) working group has now been formed which aims to lead to the formation of a working group on these issues. The group aims to develop one or more congestion control algorithms, associated extensions, and evaluation criteria. Furthermore it has been proposed that certain practices, such as 'circuit-breaker' conditions, to provide operational limits on congestion control algorithms, and feedback messages, may be tackled in other groups such as AVTCORE and AVTEXT respectively.

Thus there is some movement to attempt to develop new algorithms better suited to media transport, but these efforts will clearly take a considerable time to reach fruition. Whilst TFRC has some perceived issues it still provides the best existing solution for media transport.



### **[3.3](#) ECN response**

As mentioned above and in accordance to [[rfc3168](#)], the actual response to the reception of an ECN-CE marked packet MUST normally be the same as that of a lost packet. However there are a number of contexts where one may also be interested in more varied approaches. We expand on this in [Section 5](#) below.

## **[4.](#) Application Layer Congestion Response**

Whilst the congestion control algorithm may decide to alter the rate at which the application should operate, in the case of media applications this process is not as straightforward as the case of bulk data. The different media engines and codecs in use may only have limited adaptation ranges, thus, this limitation needs to be a consideration when adapting the rate. Furthermore the application needs to be aware of the capability of the specific codecs in terms of their ability to switch configuration mid-stream (without loss of fidelity), which may impose further limits on the modes of operation.

One approach for achieving a lower generation of data is through reduced sampling of the media (e.g., voice or video). In the case of video, this may also involve slower frame rates. Specific recommendations that describe how applications should respond to congestion in the context of supporting the algorithmic characteristics of a congestion control algorithm are outside the scope of this document.

## **[5.](#) Other Reactions**

In addition to the activation of congestion control algorithm, other reactions can be used or leveraged by an application in response to CE. We divide these other potential reactions into two categories: signaling and fault tolerance. We note that these other reactions are considered symmetric because they require downstream peer support. We also point out that activation of other reactions represents an example of an on-demand and as-needed approach in responding to CE.

### **[5.1](#) Signaling**

#### **[5.1.1](#) RSVP**

The resource Reservation Protocol (RSVP) can be used to signal a desired set of path characteristics (e.g., bandwidth, delay) in response to CE feedback [[rfc2205](#)]. Its operation is based on the use of PATH messages sent downstream hop-by-hop from the source to a destination that specify requested forwarding characteristics. In return, the destination sends a hop-by-hop RESV message upstream towards the source confirming the resources that have been reserved for that flow.



[rfc3181] defines a priority policy element that specifies both an allocation and defending priority. This dual specification supports the use of preemption of existing reservations. [draft-priority-rsvp] is a work-in-progress that defines a new policy element that only conveys priority during reservation establishment. This latter effort also presents several reservation models, including one that describes engineered resources set aside for priority users.

### **5.1.2 Differentiated Services**

Unlike RSVP and its use of a separate signaling mechanism to reserve resources, Differentiated Services (diff-serv) uses code points within the IP header to convey the forwarding behavior of that packet [rfc2474]. This may range from various drop precedence values to a code point that signifies low delay and low loss (i.e., characteristics attributed to real time flows).

As in the case of RSVP, applications could rely on the reception of CE feedback to initiate a subsequent setting of diff-serv code points to provide additional protection or explicit association of forwarding characteristics of a given flow of packets. In addition, the setting of diff-serv code points would be done on an as-needed basis in reaction to CE feedback. Recommendations concerning specific diff-serv values are outside the scope of this document.

## **5.2 Fault Tolerance**

Fault tolerance is another category of reactions that may be used by applications in response to CE feedback. In some cases, these efforts may contribute to an increase in traffic load in order to add protection and resiliency to a flow.

**Redundant Transmissions:** This approach is based on a source sending duplicate payloads that can be used to compensate for lost packets. Given that ECN marks the packet and forwards it towards the destination (instead of dropping it), this approach can be considered extreme in terms of being network unfriendly. Its positive value may emerge in cases where a path has several downstream congestion points. However, its actions of producing redundant packets still associates a high measure of greedy use of resources.

**Application Layer Forward Error Correction (FEC):** This approach also adds additional overhead to the flow in order to compensate for potential packet loss. And as the case of redundant transmissions, the value of this approach is probably better realized when there exists multiple downstream congestion points. However, the impact of the overhead is minimized by having one (or a few) additional packet(s) used to compensate for the loss of a set of packets.



Codec Swapping: This approach involves changing codecs to either reduce load or achieve an improvement in compensating for lost packets.

### 5.3 Alternative Reaction for Emergency Communications

As mentioned in [rtp-ecn], the default reaction on the reception of these ECN-CE marked packets MUST be to provide the congestion control algorithm with a congestion notification that triggers the algorithm to react as if packet loss had occurred. There MAY be an alternative reaction if it is considered safe for deployment. An example of the need for an alternative reaction would be the case of Emergency Telecommunications Service (ETS) [rfc3689, [rfc4190](#)], where an improvement in QoS or a higher probability of session establishment and forwarding of traffic is of high interest.

It is proposed that certain authorized ETS flows may be permitted to employ either a substantially less aggressive back-off algorithm than the default algorithm, or some level of exemption from reacting to ECN marked packets. This alternative reaction will benefit these flows as the marks would normally be considered as equivalent to lost packets, which would effectively increase the loss level, which in turn will generally result in the reduction of flow rate. This applies to all flows that utilize some form of the rate control that is inversely proportional to the loss rate, which includes TCP-like algorithms or equation-based approaches.

Simulations of the use of ECN exemption with TFRC and have found that it has limited effect on the normal flows with low numbers of exempt flows. A half-dumbbell network was used with a RED router queue configured using the settings recommended by Sally Floyd. The candidate flows are 1Mbit/s each with a backhaul 100Mbit/s link. In the standard case where 1% of flows would be exempt the remaining flows achieve 99.99% of the bandwidth that they would achieve without the presence of the exempt flows. This is what would be expected from the simple calculation of the allocation, given that the exempt flows achieve their full rate (1Mbit/s); With 100 normal plus 1 exempt flow, assuming that the exempt flow uses 1Mbit/s, the remaining capacity is 99Mbit/s which is divided between the 100 normal flows. Whilst when 101 normal flows are run over the 100Mbit/s link they would have to share it evenly, so it works out thus:  $((99/100)/(100/101))*100=99.99\%$ . In the case of 5% exempt flows then the proportion is very slightly lower at  $((95/100)/(100/105))*100=99.75\%$ . Both these calculations are borne out in the simulation runs.

The level of exemption employed can be altered in a number of ways. Two simple approaches would be to either set a threshold number of ECN marked packets that could be considered as a loss, and another approach would be to set a percentage threshold of ECN marked packet that would





be considered as a loss.

It should be noted that in the simulations the end-to-end delay of the packets within the flows was monitored and the relative delay of the exempt flows apparently rises somewhat when exemption is enacted. However what is actually occurring is that the 'normal' flows are reducing their throughput and are thus reducing their latency somewhat. There is normally some limited latency when using loss-based techniques such as TFRC because it fills the queues to ascertain the link capacity and maintains that level of delay throughout a session. However the level of latency is clearly limited by the queue sizes in the network and on media specific links these queue sizes are typically quite small, so the resulting latency is limited.

Furthermore in the case where media flows employing TFRC, or any other congestion control algorithm (e.g. delay-based), are sharing a bottleneck link with TCP flows then the queues will be filled by the TCP flows and the latency will be kept near or at a their maximum despite any other flows.

## **6. IANA Considerations**

This document requires no actions from IANA.

## **7. Security Considerations**

The reliance on accurate and un-modified RTCP information means that SRTP needs to be used, or any other mechanism that helps prevent modification of RTCP feedback packets.

## **8. Acknowledgements**

TBD

## **9. References**

### **9.1 Normative**

- [rfc2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [rfc2205] Braden, B., et. al., "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", [RFC 2205](#), September 1997
- [rfc2209] Braden, R., L. Zhang, "Resource Reservation Protocol (RSVP) -- Version 1 Message Processing Rules", [RFC2209](#)



September 1997

- [rfc2474] Nichols, K., et. al., "Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers", [RFC 2474](#), December 1998
- [rfc3168] Ramakrishnan, K., et. al., "The Addition of Explicit Congestion Notification (ECN) to IP", [RFC 3168](#), September, 2001
- [rfc3181] Herzog, S., "Signaled Preemption Priority Policy Element", [RFC 3181](#), October 2001
- [rfc3448] Handley, M., et. al., "TCP Friendly Rate Control (TFRC): Protocol Specification", [RFC 3448](#), January 2003
- [rfc4867] Sjöberg, J., et. al., "RTP Payload Format and File Storage Format for the AMR and AMR-WB Audio Codecs", [RFC 4867](#), April 2007
- [rfc5104] Wenger, S., et. al., "Codec Control Messages in the RTP Audio-Visual Profile with Feedback (AVPF)", [RFC 5104](#), February 2008
- [rfc6679] Westerlund, M., et. al., "Explicit Congestion Notification (ECN) for RTP over UDP", [RFC 6679](#), IETF, July 2012

## **9.2 Informative**

- [[draft-rtp-tfrc](#)] Gharai, L., C. Perkins, "RTP with TCP Friendly Rate Control", work-in-progress, Sept 2011
- [Goog1] [http://code.google.com/apis/talk/call\\_signaling.html](http://code.google.com/apis/talk/call_signaling.html)
- [tr26.114] "IMS; Multimedia telephony; Media Handling and Interaction", 3GPP, version 10, April 2011
- [ts36.300] "E-UTRA and E-UTRAN Overall Description, Stage 2", 3GPP, Release 10, September, 2011
- [RFC4340] Kohler, E., et. al, Datagram Congestion Control Protocol (DCCP), [RFC4340](#), March 2006
- [RFC4342] Floyd, S., et. al., "Profile for DCCP Congestion Control ID 3: TFRC", [RFC 4342](#), March 2006
- [RFC4828] Floyd, S., E. Kohler, "TFRC: The Small Packet



Variant", [RFC 4828](#), April 2007

[rfc3689] Carlberg, K., Atkinson, R., "General Requirements for  
Emergency Telecommunications Service (ETS)", [RFC 3689](#),  
February 2004

[rfc4190] Carlberg, K. et, al., "Framework for Supporting  
Emergency Telecommunications Service (ETS) in  
IP Telephony", [RFC 4190](#), November 2005

[rfc3714] Floyd, S., Kempf, J., "IAB Concerns Regarding Congestion  
Control for Voice Traffic in the Internet", [RFC 3714](#),  
March 2004

[bcp145] Eggert, L., Fairhurst, G., "Unicast UDP Usage Guidelines  
for Application Designers", [RFC 5405](#), [BCP 145](#), November 2008

[ITU.G114.2003]  
International Telecommunications Union, "One-way  
transmission time", ITU-T Recommendation G.707, May 2003.

#### Author's Addresses

Piers O'Hanlon  
University of Oxford  
Oxford Internet Institute  
1 St Giles  
Oxford OX1 3JS  
United Kingdom

Email: [piers.ohanlon@oii.ox.ac.uk](mailto:piers.ohanlon@oii.ox.ac.uk)

Ken Carlberg  
G11  
1600 Clarendon Blvd  
Arlington VA  
USA

Email: [carlberg@g11.org.uk](mailto:carlberg@g11.org.uk)

