IETF D. Kutscher
Internet-Draft NEC
Intended status: Informational M. Kuehlewind
Expires: December 25, 2012 IKR University of Stuttgart
June 23, 2012

Towards Adaptive Congestion Management for Interactive Real-Time Communications
draft-kutscher-rtc-conex

Abstract

Real-time communication frameworks such as RTCWEB can facilitate session setup for users and applications and will potentially lead to an increase of real-time traffic on the Internet. Unlike operator-controlled real-time services in IMS or similar frameworks, this traffic is likely to share capacity with other best-effort traffic. If real-time communication applications do not implement congestion control or some kind of rate adaption to react to congestion in a timely manner, this can lead to bad service quality -- and affect other, more adaptive TCP-based applications and in worse case network stability. Instead of introducing operator-enforced bandwidth limitations or even risking blocking of such 'over-the-top' traffic altogether, we propose incentivizing applications and real-time transport protocols to react to congestion by applying congestion exposure (ConEx) and corresponding policing frameworks. In this paper we discuss existing IETF mechanisms that can be used and we point at gaps that should be filled by future work.

Status of this Memo

Copyright Notice

Table of Contents

1. Motivation

Historically [ref.confarch], interactive real-time communications was intended as yet another Internet application that could be used in point-to-point or multicast settings. In addition to RTP, control protocols such as SAP and SIP have been designed to announce/initiate sessions. Telephony-oriented call-admission and charging functions were not deemed necessary or left for future work. For congestion control RTCP-based statistical feedback was considered to be sufficient for most applications -- or application and protocol developers have been referred to TFRC [RFC5348] as a default solution.

The development of network/service architectures such as IMS and NGN involved the applying of SIP for call-control, call-admission and charging, based on a model where SIP-based call initiation would lead to corresponding resource reservation. Real-time communication frameworks such as RTCWEB can facilitate session setup for users and applications and will potentially lead to an increase of real-time traffic on the Internet, without necessarily using established session initiation protocols, i.e., SIP. In addition, the availability of better and affordable hardware as well as application support had made video calls (from QCIF to HDTV resolutions) much more popular. Still lacking proper congestion control, such traffic can be a threat to other, more adaptive TCP-based applications.

In this perspective, operators can perceive OTT interactive real-time communication more threatening than P2P traffic (that largely uses TCP and thus congestion control). There is a risk that operators will extend attempts to detect and limit the usage of OTT interactive real-time communication through Deep-Packet Inspection and corresponding filtering.

Instead of promoting a arms race in which RTCWEB service providers have to continuously invent new signaling and real-time traffic obfuscation mechanisms, we call for a better way: make congestion visible to the network and let applications decide how to react. Operators can police senders based on their response to congestion, the current state of the network and other criteria such as service contracts. The approach would enable both, quick response to congestion as well as long-term "fairness", and it would enable interactive real-time communication to share networks with other applications, for example other "best-effort" traffic.

In this position paper, we list some requirements for managing congestion for interactive real-time communication, we explain how the ConEx framework can be used, and we list a few mechanisms that could be used and point at gaps that need more work.

2. Requirements

Congestion control independent of signaling protocols: VoIP, RTCWB-based and P2P applications can use different signaling or management protocols, although they are often using RTP or UDP-based transport in a similar fashion. Instead of tying resource reservation and/or congestion control to those signaling protocols, there should be transport-layer mechanisms that work independent of signaling protocols.

Timely and accurate feedback: In order to be effective for mitigating network congestion, feedback must be generated and transferred to a sender fast (RTT timescale). Depending on the implemented congestion response mechanism the feedback information has to provide sufficient accuracy. In the best case the exact number of bytes that have determined the congestion is reported (congestion byte-volume). At the same time, congestion feedback should not induce a large data transmission overhead.

Simple, standardized congestion response: Interactive real-time communication can be congestion-controlled, but senders (developers of corresponding applications) have to know concretely how to respond. For instance, how shall a congestion signal be translated into a codec parameter change? Simple and standardized profiles should be provided other than TCP congestion control as interactive real-time communication require a more smooth reaction or rate adaptations in concrete steps up to a certain minimum rate. If the minmum requirement can not be provided by the network the services should be interrupted as the user will not be able to actually use the service.

Signaling protocol support to negociate rate adaption: Although the congestion control loop itself should be independent of a signaling protocol, the signaling protocol should be able to negotiate the use of rate adaptation features. This can be realized by an offer/answer process at the beginning of the communication where e.g. availability of alternative media types at both endpoints are declared. In the best case no further interaction during the communication is needed.

Incentives for senders and receivers: Unless there are good incentives for senders and receivers (in case of receiver-driven congestion control) to react to congestion signals by reducing their current congestion contribution, there can be a winning strategy in using FEC formats to increase the probability of transmitting enough media data in the presence of congestion -- which would penalize other users even more.

3. ConEx framework

Congestion exposure can help meeting the above-mentioned requirements, because it enables both endpoints and the network to see congestion and to act on it at different time-scales (RTT timescale for congestion control loops but also longer timescales for achieving some form of fairness, accounting etc.).

The ConEx abstract mechanism is defined in [I-D.ietf-conex-abstract-mech]. The key mechanism is a new IP layer congestion exposure signal that is exposes current congestion contribution per flow and can be used to calculate rest-of-path congestion (when comparing ConEx information to ECN signals).

By exposing congestion contribution on the IP layer, ConEx enables application-independent network-based reaction to congestion: In addition to network devices and transport senders and receivers ConEx assumes policy devices that can react on the exposed information. An ingress policers can thus implement operator policies for interactive real-time as well as for other traffic as described below. A policer will maintain a certain congestion buget per user. This will allow the user/application to react to congestion appropriately.

For interactive real-time communication, the most important benefit would be that traffic management can be informed about RTCWEB or other real-time media traffic and its contribution to congestion. Network operators can apply different policies, for example ensuring a useful resource distribution between RTCWEB and other applications' flows -- considering users or user groups and other criteria. Significant, possibly long-term congestion contribution can -- depending on the type of network and the service agreement -- lead to stricter management (QoS class reduction, packet dropping, charging).

For RTP senders, the notion of congestion-aware traffic management in the network can provide useful incentives to really react to congestion signals fast and adequately. Senders can see the current congestion contribution of individual flows, sets of flows representing an application or user, and can decide how to split the available budget among those flows.

4. Mechanisms

We now list and assess a few existing IETF mechanisms that should be used -- and we point at some gaps that should be filled by future work.

4.1. RTP Control Protocol Extended Reports (RTCP XR)

RFC 3611 [RFC3611] defines the Extended Report (XR) packet type for RTCP and a set of specific extensions such as detailed packet loss/ duplication reports, wallclock timestamp reports, statistics summary, and VoIP metric reports.

RFC 3611 provides a general RTCP report format extension that can be used as a basis for additional reports, however it does not change the general RTCP specification [RFC3550] on reporting intervals (5 seconds minimum interval). For congestion exposure aiming to support congestion control, a much shorter interval would be needed.

4.2. Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)

The Extended RTP Profile RTCP-based Feedback [RFC4585] specifies additional feedback modes, including an immediate feedback mode (originally intended for triggering re-transmissions). Immediate feedback is usually intended for point-to-point and smaller group communication sessions. RFC 4585 defines different types of feedback messages: 1) Transport layer messages, 2) Payload-specific messages, and 3) Application layer FB messages.

4.3. Explicit Congestion Notification (ECN) for RTP over UDP

ECN for RTP/URP [I-D.ietf-avtcore-ecn-for-rtp] defines an RTCP Extended Report (XR) block for periodic ECN feedback and a new RTCP transport feedback message for timely reporting of congestion events, using RFC 4585’s immediate reporting mechanism.

The ECN feedback report format provides, amongst other information, a list of absolute numbers of received ETC(0)/ETC(1)/ECN-CE/not-ECT packets from a specified SSRC.

RTP receivers are expected to return ECN RTCP feedback as soon as possible using this report format. An RTP sender, receiving such feedback, is expected to provide its congestion control algorithm with a congestion notification. In a sender-driven congestion control scheme, rate adaptation, employing TFRC [RFC5348] could be used. There are also receiver-driven congestion control mechanisms (e.g., employing layered coding), where the RTP receiver can react directly, e.g., by unsubscribing from receiving an encoding layer stream, however, this is normally to happen on larger time-scales, i.e, not as an immediate reaction to congestion notification.

The specification does currently not include accurate congestion feedback, i.e., volume-based feedback.

4.4. Circuit Breakers for RTP Unicast Sessions

RTP circuit breakers [I-D.perkins-avtcore-rtp-circuit-breakers] are intended to serve as a last-resort reaction to network congestion (lacking proper congestion control mechanisms), i.e., they are conditions under which an RTP flow should cease to transmit media to protect the network from excessive congestion as the service experience will be very low anyway.

Congestion can be determined by using regular RTCP reporting or by employing the more immediate and detailed mechanisms as described above.

4.5. Missing Elements

4.5.1. Congestion Feedback Mechanisms

ConEx policing and auditing functions are expected to operate on congestion-volume, so it can be helpful to make the congestion feedback and declaration work on the basis of congestion volume instead of (marked) packet counters. ECN for RTP already enables quite accurate feedback (more than one congestion event per RTT through its ECN counters). For constant packet size payload formats this information can be directly translated to congestion volume. For other payload format, such as video formats this is not the case, so it could be useful to extend congestion reports by byte volume counters.

4.5.2. Sender-driven Congestion Control Algorithms

Clearly, what is needed is a viable congestion control scheme that enables senders to re-act to congestion notification appropriately. It is not sufficient to refer to TFRC and leave this to application developers. We cannot specify a complete solution in this memo, but would like to outline a few key elements.

Rate adaptation: There should be a defined downgrading path for senders along the lines of TRFC, i.e., providing a smoother reaction than AIMD. Such rate adaptation downgrading paths could be specified in payload type definitions.

RTP payload type adaptation: SDP and its usage in SIP already allows for declaring/negotiating a set of supported/acceptable payload types (e.g., video codec with format parameter definition) for a certain media type. A set of well-known or standardized payload types should be documented that allows for an appropriate rate adaptation.

Session description/negotiation extensions: Establishing a rate/payload type adaptation scheme required a common understanding of senders and receivers about the supported variants. Session initiation mechanisms (SIP, offer/answer, SDP) might need extensions in this direction.

4.5.3. Incentivizing receiver-driven congestion control

When layered video coding becomes more commonplace, there is still the question how to motivate receivers to make use of it. We are proposing a ConEx-based policing framework for this. Again, in order to arrive at reasonable layer selection decisions, receivers would need accurate information of current congestion contribution and information that enables them to (de-)select layers proportionally to the intended rate reduction.

4.5.4. Mechanisms for Congestion Exposure and Policing

RTP-friendly policing regimes: The rate adaptation for RTP flows is not guaranteed to be as responsive as TCP congestion control can be. For example, it can take a while until codec format parameter configuration changes have an effect on the bytes to be sent. For more drastic changes, e.g., changing codecs, the delay can be several seconds. Congestion budgets at policers should allow for such grace periods. Note that this is a different concept than the credit concept for audit functions in [I-D.ietf-conex-abstract-mech].

operator policies: In general, ConEx allows for application-independent policing and enables operators to implement different policing approaches. For example, in networks with stricter per-user accounting continued congestion contribution for traffic from the user and for traffic to the user can lead to either charging or service downgrading.

sender stratgies: Knowing such policies senders can decide how to respond to congestion signals. On a host, congestion contribution for a number of flows (that share a path) can be considered together, and applications can choose which flow to regulate etc. For example, in a multimedia conference call, it can be OK to adapt the RTP video flow only.

5. Summary

Summarizing, interactive real-time communication and its expected usage in RTCWEB applications requires solid congestion control to ensure the success of RTCWEB and good user experience in operated networks. Congestion exposure is a useful mechanism to make congestion contribution visible in the network, so that senders/ receivers can react adequately and traffic management can implement operators policies. The policing element can incentivize applications to really make use of available rate reduction mechanisms. As a ConEx policer allows only a certain congestion buget per user, in persitent congestion situation a real-time application should increase the traffic volume by e.g. introducing more redundancy, but reaction with an adequate rate control.

Existing elements such as immediate RTCP feedback and ECN for RTP can be used, but there are also some gaps such as standardized adaptation mechanisms and the mentioned signaling support. For making congestion exposure work effectively, policing and auditing elements should be prepared to give more congestion credit to real-time streams, and useful operator policies for real-time traffic should be defined. The IETF should work on at least some of these topics. The presentation at the workshop will discuss ideas for addressing those gaps in more detail.

6. Security Considerations

There certainly are :-)

7. Informative References

[I-D.ietf-avtcore-ecn-for-rtp]
Westerlund, M., Johansson, I., Perkins, C., O'Hanlon, P., and K. Carlberg, "Explicit Congestion Notification (ECN) for RTP over UDP", draft-ietf-avtcore-ecn-for-rtp-08 (work in progress), May 2012.

[I-D.ietf-conex-abstract-mech]
Mathis, M. and B. Briscoe, "Congestion Exposure (ConEx) Concepts and Abstract Mechanism", draft-ietf-conex-abstract-mech-04 (work in progress), March 2012.

[I-D.perkins-avtcore-rtp-circuit-breakers]
Perkins, C. and V. Singh, "RTP Congestion Control: Circuit Breakers for Unicast Sessions", draft-perkins-avtcore-rtp-circuit-breakers-00 (work in progress), March 2012.

[RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time

Applications", STD 64, RFC 3550, July 2003.

[RFC3611] Friedman, T., Caceres, R., and A. Clark, "RTP Control Protocol Extended Reports (RTCP XR)", RFC 3611, November 2003.

[RFC4585] Ott, J., Wenger, S., Sato, N., Burmeister, C., and J. Rey, "Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)", RFC 4585, July 2006.

[RFC5348] Floyd, S., Handley, M., Padhye, J., and J. Widmer, "TCP Friendly Rate Control (TFRC): Protocol Specification", RFC 5348, September 2008.

[ref.confarch] Handley, M., Crowcroft, J., Bormann, C., and J. Ott, "The Internet Multimedia Conferencing Architecture", draft-ietf-mmusic-confarch-03 (work in progress), July 2000.

(Expired)

Authors' Addresses

Dirk Kutscher
NEC
Kurfuersten-Anlage 36
Heidelberg,
Germany

Phone:
Email: kutscher@neclab.eu

Mirja Kuehlewind
IKR University of Stuttgart
Germany

Email: mirja.kuehlewind@ikr.uni-stuttgart.de