University of Glasgow

RTP Congestion Control Feedback

Colin Perkins

Reporting on draft-ietf-avtcore-cc-feedback-message-09, that's is co-authored with Zahed Sarker, Varun Singh, and Michael A. Ramalho

Draft Status

- draft-ietf-avtcore-cc-feedback-message-08 went to IESG review
- Comments from TSVART, GenART, IANA, Benjamin Kaduk, Éric Vyncke, Erik Kline, Magnus Westerlund, Martin Duke, and Robert Wilton
- Mostly minor editorial nits – a small number of technical changes to review

Editorial Changes

- Add to Abstract: “An effective RTP congestion control algorithm requires more fine-grained feedback on packet loss, timing, and ECN marks than is provided by the standard RTP Control Protocol (RTCP) Sender Report (SR) and Receiver Report (RR) packets”
- Don’t reference the terminology defined in RFCs 3551 and 3611, since those RFCs don’t define terminology
- Clarify: “RTCP packets do not contain a sequence number, so loss of feedback packets has to be inferred based on the time since the last feedback packet.”
- Clarify that it’s the media sender that should reduce its rate if multiple feedback packets are lost
- Since draft-alvestrand-rmcat-remb expired long ago, reference “Mechanisms that convey the receiver's estimate of the maximum available bit-rate” rather than that explicit draft

Packet Format Description

The contents of each 16-bit packet metric block comprises the L, ECN, and ATO fields as follows:

o L (1 bit): is a boolean to indicate if the packet was received. 0 represents that the packet was not yet received and all the subsequent bits (ECN and ATO) are also set to 0. 1 represents that the packet was received and the subsequent bits in the block need to be parsed.

The contents of each 16-bit packet metric block comprises the R, ECN, and ATO fields as follows:

o Received (R, 1 bit): is a boolean to indicate if the packet was received. 0 represents that the packet was not yet received and the subsequent 15-bits (ECN and ATO) in this 16-bit packet metric block are also set to 0 and MUST be ignored. 1 represents that the packet was received and the subsequent bits in the block need to be parsed.

- Rename field L → R to match usage
- Clarify that if the packet is lost, the bits indicating its arrival time and ECN mark are set to zero and MUST be ignored

Relation to RFC 6679 ECN Feedback

The RTCP ECN Feedback Packet is not useful when ECN is used with the RTP Congestion Control Feedback Packet defined in this memo since it provides duplicate information. Accordingly, when congestion control feedback is to be used with RTP and ECN, the SDP offer generated MUST include an "a=ecn-capable-rtp:" attribute to negotiate ECN support, along with an "a=rtcp-fb:" attribute with the "ack" parameter "ccfb" to indicate that the RTP Congestion Control Feedback Packet is to be used for feedback. The "a=rtcp-fb:" attribute MUST NOT include the "nack" parameter "ecn", so the RTCP ECN Feedback Packet will not be used.

The RTCP ECN Feedback Packet is not useful when ECN is used with the RTP Congestion Control Feedback Packet defined in this memo since it provides duplicate information. When congestion control feedback is to be used with RTP and ECN, the SDP offer generated MUST include an "a=ecn-capable-rtp:" attribute to negotiate ECN support, along with an "a=rtcp-fb:" attribute with the "ack" parameter "ccfb" to indicate that the RTP Congestion Control Feedback Packet can be used. The "a=rtcp-fb:" attribute MAY also include the "nack" parameter "ecn", to indicate that the RTCP ECN Feedback Packet is also supported. If an SDP offer signals support for both RTP Congestion Control Feedback Packets and the RTCP ECN Feedback Packet, the answering party SHOULD signal support for one, but not both, formats in its SDP answer to avoid sending duplicate feedback.

When using ECN with RTP, the guidelines in Section 7.2 of [RFC6679] MUST be followed to initiate the use of ECN in an RTP session. The guidelines in Section 7.3 of [RFC6679] MUST also be followed about ongoing use of ECN within an RTP session, with the exception that feedback is sent using the RTCP Congestion Control Feedback Packets described in this memo rather than using RTP ECN Feedback Packets. Similarly, the guidance in Section 7.4 of [RFC6679] around detecting failures MUST be followed, with the exception that the necessary information is retrieved from the RTCP Congestion Control Feedback Packets rather than from RTP ECN Feedback Packets.

- Change ECN feedback signalling, to allow more graceful fallback to RFC 6679 format
- Mandate that the guidelines in RFC 6679 for checking that ECN actually works on the path need to be followed

Clarify Security Considerations

```
11.  Security Considerations

   The security considerations of the RTP specification [RFC3550], the
   applicable RTP profile (e.g., [RFC3551], [RFC3711], or [RFC4585]),
   and the RTP congestion control algorithm that is in use (e.g.,
   [RFC8698], [RFC8298], [I-D.ietf-rmcat-gcc], or [RFC8382]) apply.

   A receiver that intentionally generates inaccurate RTCP congestion
   control feedback reports might be able trick the sender into sending
   at a greater rate than the path can support, thereby causing
   congestion on the path.  This will negatively impact the quality of
   experience of that receiver.  Since RTP is an unreliable transport, a
   sender can intentionally leave a gap in the RTP sequence number space
   without causing harm, to check that the receiver is correctly
   reporting losses.

   An on-path attacker that can modify RTCP congestion control feedback
   packets can change the reports to trick the sender into sending at
   either an excessively high or excessively low rate, leading to denial
   of service.  The secure RTCP profile [RFC3711] can be used to
   authenticate RTCP packets to protect against this attack.
```

```
11.  Security Considerations

   The security considerations of the RTP specification [RFC3550], the
   applicable RTP profile (e.g., [RFC3551], [RFC3711], or [RFC4585]),
   and the RTP congestion control algorithm that is in use (e.g.,
   [RFC8698], [RFC8298], [I-D.ietf-rmcat-gcc], or [RFC8382]) apply.

   A receiver that intentionally generates inaccurate RTCP congestion
   control feedback reports might be able trick the sender into sending
   at a greater rate than the path can support, thereby causing
   congestion on the path.  This will negatively impact the quality of
   experience of that receiver, and potentially cause denial of service
   to other traffic sharing the path and excessive resource usage at the
   media sender.  Since RTP is an unreliable transport, a sender can
   intentionally drop a packet, leaving a gap in the RTP sequence number
   space without causing serious harm, to check that the receiver is
   correctly reporting losses (this needs to be done with care and some
   awareness of the media data being sent, to limit impact on the user
   experience).

   An on-path attacker that can modify RTCP congestion control feedback
   packets can change the reports to trick the sender into sending at
   either an excessively high or excessively low rate, leading to denial
   of service.  The secure RTCP profile [RFC3711] can be used to
   authenticate RTCP packets to protect against this attack.

   An off-patch attacker that can spoof RTCP congestion control feedback
   packets can similarly trick a sender into sending at an incorrect
   rate, leading to denial of service.  This attack is difficult, since
   the attacker needs to guess the SSRC and sequence number in addition
   to the destination transport address.  As with on-patch attacks, the
   secure RTCP profile [RFC3711] can be used to authenticate RTCP
   packets to protect against this attack.
```

Next Steps

- Is the WG okay with these changes?

- If yes, IESG approves and this goes to RFC Editor
- If no, need update and new review cycle